Tivoli software

IBM

Deployment Guide Series: IBM Tivoli Storage Manager V5.5

Plan for IBM Tivoli Storage Manager deployment

Perform installation and configuration

Learn from example scenarios

Budi Darmawan
Maria Coronel
Alexandre Guedes
Klavs Kabell

Redbooks

ibm.com/redbooks

International Technical Support Organization

Deployment Guide Series: IBM Tivoli Storage Manager V5.5

October 2008

SG24-7379-01

Note: Before using this information and the product it supports, read the information in “Notices” on page ix.

Second edition (October 2008)

This edition applies to Version 5, Release 5 of IBM Tivoli Storage Manager.

Contents

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®
Database 2™
DB2 Universal Database™
DB2®
Domino®
DPI®
DS4000™
DS6000™
DS8000™
Enterprise Storage Server®
FlashCopy®
General Parallel File System™
GPFS™
HACMP™
IBM Virtual Innovation Center™
IBM®
Informix®
iSeries®
Lotus®
Magstar®
MVS™
NetView®
OS/390®
OS/400®
PartnerWorld®
Passport Advantage®
POWER™
POWER6™
pSeries®
Redbooks®
Redbooks (logo) ®
SANergy®
SysBack™
System p®
System Storage™
System x™
Tivoli Enterprise Console®
Tivoli®
TotalStorage®
Virtualization Engine™
WebSphere®
z/OS®
zSeries®

The following terms are trademarks of other companies:

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

Snapshot, Network Appliance, NearStore, Data ONTAP, and the NetApp logo are trademarks or registered trademarks of NetApp, Inc. in the U.S. and other countries.

Novell, SUSE, the Novell logo, and the N logo are registered trademarks of Novell, Inc. in the United States and other countries.

Oracle, JD Edwards, PeopleSoft, Siebel, and TopLink are registered trademarks of Oracle Corporation and/or its affiliates.

mySAP, SAP R/3, SAP, and SAP logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries.

VMware, the VMware "boxes" logo and design are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions.

Java, JRE, Solaris, Sun, ZFS, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Access, Active Directory, Microsoft, SharePoint, SQL Server, Windows NT, Windows Server, Windows Vista, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Preface

The objective of this IBM® Redbooks® publication is to provide comprehensive instructions for deploying the IBM Tivoli® Storage Manager to various environments.

Readers should possess general knowledge of communication network architecture and design, basic size considerations of the Tivoli Storage Manager database, and basic pool management of Tivoli Storage Manager servers. This document is intended to be read and used by presales systems engineers and services personnel to build a customized deployment of the Tivoli Storage Manager. We expect such individuals have significant knowledge of Tivoli Storage Manager, and ideally the reader has attended Tivoli Storage Manager basic and advanced training classes.

The reader must be familiar with the following topics:

- Storage management concepts
- Network topologies
- Distributed systems architectures and configuration

This book is a valuable addition to, and can be read in conjunction with the existing product documentation.

The team that wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, Austin Center.

Budi Darmawan is a Project Leader at the International Technical Support Organization, Austin Center. He writes extensively and teaches IBM classes worldwide on all areas of Tivoli systems management. Before joining the ITSO 9 years ago, Budi worked in IBM Global Services Indonesia as lead implementer and solution architect. Budi is also an ITIL® Certified Service Manager. His current interests are general systems management, availability and capacity management, operational automation, and Java™ programming.

Maria Coronel is an Advisory IT Specialist at the Tivoli Support team in Sydney, Australia. She has 14 years of experience in the storage field. Before joining the services team in 1999, Maria worked as a systems administrator responsible for the organization's backup and recovery infrastructure. Maria joined the Tivoli

Support team in 2002, specializing in the Tivoli Storage Manager products suite. She holds a master's degree in management from University of Sydney. Her broad experience with the Tivoli Storage Manager product suite, as a client and in implementing the product during her years in the services team, make her a valuable member of the Tivoli Storage Manager Support team.

Alexandre Guedes is an IT Specialist for MTS Brazil, working with the TotalStorage® Productivity Center and the Tivoli Storage Manager team, implementing and supporting Tivoli Storage Manager solutions and all their complementary products. He has been working at IBM for 10 years. He has 4 years of experience with Tivoli Storage Manager, and his areas of expertise include mobile support, network administration in a Microsoft® Windows® environment, IBM Lotus® Domino® support and administration, and IBM DB2® support. He is an IBM Certified Deployment Professional, an IBM Certified Storage Administrator, and a Microsoft Certified Professional.

Klavs Kabell is a Senior System Consultant with IT-WIT Aps in Denmark, where he has been supporting and designing Tivoli Storage Manager solutions for 7 years. His main areas of expertise are designing, implementing, and migrating large Tivoli Storage Manager environments. He has broad experience with running Tivoli Storage Manager on IBM AIX®, HP-UX, and Microsoft Windows. His complimentary product experience includes Data Protection for SAP®, Oracle®, SQL, and SharePoint®. He is an IBM Certified Instructor in Tivoli Storage Manager V5.5 and teaches classes in Europe. He is also a certified Tivoli Storage Manager Deployment Professional V5.5 and a certified Tivoli Storage Manager Storage Administrator V5.

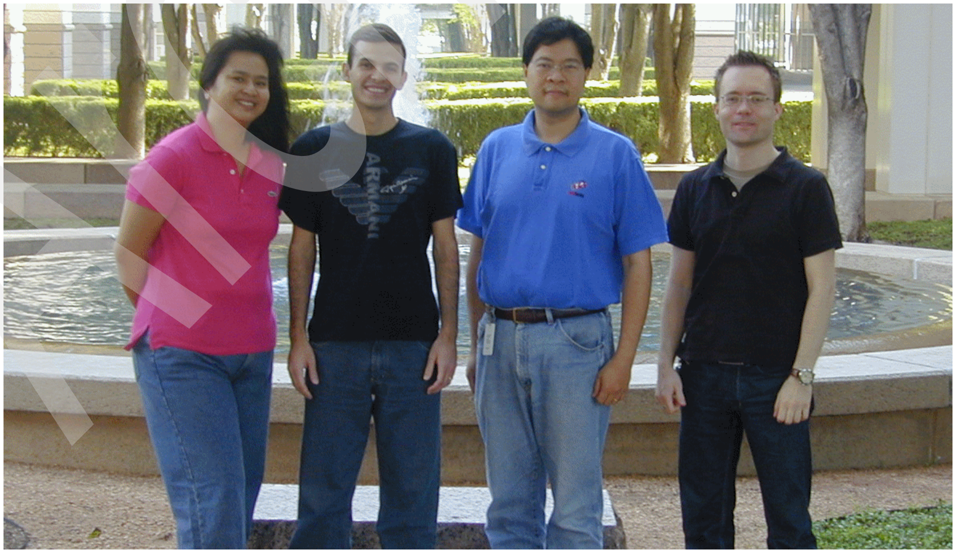

Figure 1   Maria Coronel, Alexandre Guedes, Budi Darmawan, Klavs Kabell

Thanks to the following people for their contributions to this project:

Nancy Crumpton
International Technical Support Organization, Austin Center

Terry Copeland
IBM Software Group, Tivoli Systems

Thanks to the authors of the first edition of this book, *Deployment Guide Series: IBM Tivoli Storage Manager Version 5.4*, SG24-7379, published in March 2007: Alexandre Guedes, Chris Hildebrand, Klavs Kabell, Marcia Kubo, and Bart Jacob.

Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

Your efforts will help increase product acceptance and client satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  ibm.com/redbooks

- Send your comments in an e-mail to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Summary of changes

This section describes the technical changes made in this edition of this IBM Redbooks publication. This edition may also include minor corrections and editorial changes that are not identified.

October 2008, second edition

This revision reflects the addition, deletion, or modification of new and changed information.

New information

- This edition is updated to cover the new features in Tivoli Storage Manager V5.5
- This edition adds chapters on backup recommendations and best practices for VMware®, Microsoft SQL Server®, Network Data Management Protocol (NDMP), Microsoft SharePoint, and Tivoli Storage Manager security features.

Part 1

Planning

This part covers various topics that assist the reader in planning the deployment of IBM Tivoli Storage Manager Version 5.5. It includes the following chapters:

1

Solution introduction

In this chapter we provide an overview of Tivoli Storage Manager concepts. This includes a high-level technical introduction to Tivoli Storage Manager, its architecture, and basic concepts. We provide an overview of the new functionality and changes for the latest version (V5.5).

We discuss the following topics:

1.1 IBM Tivoli Storage Manager overview

Tivoli Storage Manager is the premium product of choice for many organizations to achieve an efficient and effective enterprise-wide storage solution. It provides solutions for data protection, archiving, disaster recovery planning, space management, database and application protection, bare machine recovery, and data retention. Tivoli Storage Manager supports more than 44 operating platforms, using a consistent graphical user interface.

Tivoli Storage Manager provides:

- Centralized administration for data and storage management
- Fully automated data protection
- Efficient management of information growth
- High-speed automated server recovery
- Full compatibility with hundreds of storage devices, and local area network (LAN), wide area network (WAN), and storage area network (SAN) infrastructures
- Optional customized backup solutions for major groupware, enterprise resource planning (ERP) applications, and database products

Tivoli Storage Manager is the premier choice for complete storage management in mixed platform environments. Storing, protecting, and managing data growth are now among the major challenges of businesses today, requiring solutions that go beyond traditional backup and recovery methods.

Information has become the key asset of organizations and one of its most important competitive differentiating factors. Temporary inaccessibility or the complete loss of data have a huge financial impact and can drive organizations out of business. The inability to manage data has a negative impact on an organization's profitability and limits its ability to grow.

Tivoli Storage Manager protects an organization's data against hardware failures and other errors by storing backup and archive copies of data in offline storage. It can be scaled to protect hundreds of computers ranging from mobile computers to mainframes, running a variety of different operating systems connected through the Internet, WANs, LANs, or SANs. Centralized Web-based management, smart data move-and-store techniques, and comprehensive policy-based automation work together to minimize data protection administration costs and the impact on both computers and networks.

The base functions provided by Tivoli Storage Manager and its complementary products are as follows:

- Data protection (periodic backup and restore as well as disaster recovery):
  - In operational backup and data restoration, the backup process creates a copy of the data to protect against the operational loss or destruction of file or application data. The client defines how often to back up (frequency) and how many copies (versions) to retain.

    The restore process places the backup copy of data back into a client-designated system or workstation.
  - Disaster recovery refers to all the activities required to organize, manage, and automate the recovery process from a major loss of IT infrastructure and data across the enterprise. It includes processes to move data off-site into a secure vault location, to rebuild IT infrastructure, and to reload data successfully in an acceptable time frame.
- Data resource management (vital record retention, archive, and retrieval):
  - The archive process creates a copy of a file or a set of files representing an endpoint of a process for long-term storage. Files can remain on the local storage media or can be deleted. The client controls how long (using a retention period) an archive copy is to be retained.
  - The retrieval process locates the copies within the archival storage and places them back into a client-designated system or workstation.
- Space management (or hierarchical storage management (HSM)):

  This process provides the automatic and transparent movement of operational data that is not being accessed regularly from the user system disk space to a main storage repository. If the user accesses this data, it is dynamically and transparently restored to client storage.

The solution is network based, meaning that these functions are available to the entire network environment. Administration costs are minimized by centralized management of all Tivoli Storage Manager components.

Tivoli Storage Manager is available in three editions:

- Express
- Basic Edition
- Extended Edition

1.2 IBM Tivoli Storage Manager architecture

Tivoli Storage Manager is a client/server architecture as shown in Figure 1-1. Tivoli Storage Manager clients are the workstations, file servers, mobile computers, and other machines containing data that must be protected. Tivoli Storage Manager Client software is installed on these systems.

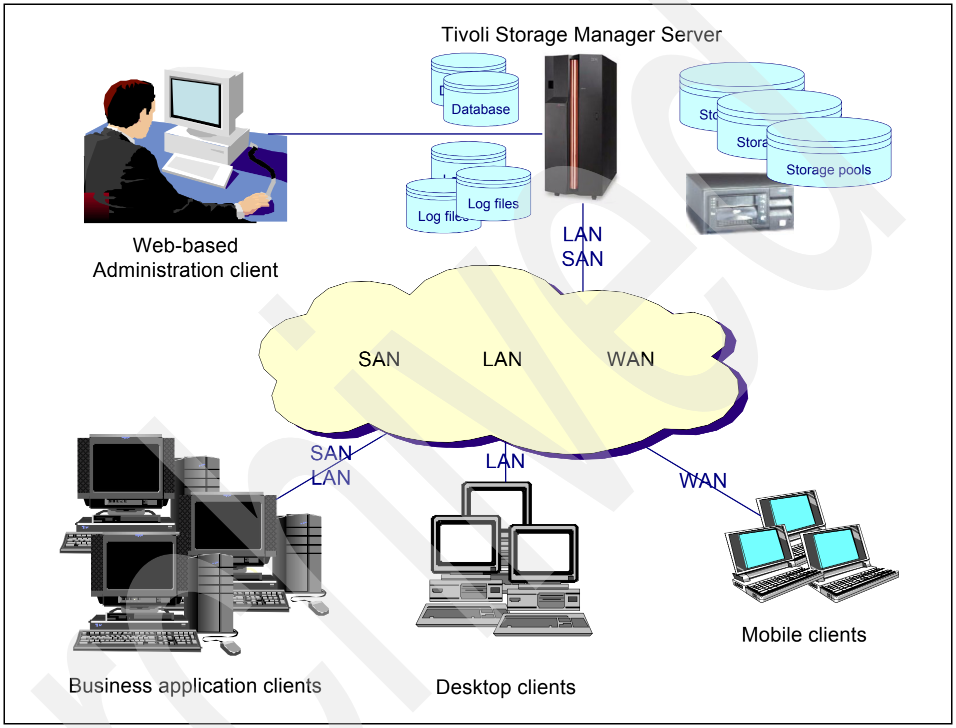

Figure 1-1 Tivoli Storage Manager architecture

Tivoli Storage Manager is based on a relational database and transaction log. The database and transaction log track metadata, what is backed up, where it is stored, and policies, schedules, and administrators.

The transaction log enables a two-phase commit, which protects the integrity of the database and allows interrupted backups and restores to be restarted. The relational database empowers Tivoli Storage Manager to perform tasks that are not possible when you use a flat file master catalog to track metadata. For example, the relational database can:

- Move data from one type of storage pool to another
- Retroactively update backed-up data when a policy changes
- Track individual files

- Schedule any type of client or administrative process
- Reclaim expired or unusable space on tapes

The Tivoli Storage Manager Client sends its data to the Tivoli Storage Manager server either by the LAN or SAN. Most backups occur through schedules, but clients can perform on-demand backups whenever they want. Clients can also perform their own restores. Tivoli Storage Manager has a new Administration Center introduced in V5.3. You can install the IBM Administration Center either on the same machine as the Tivoli Storage Manager Server or on a separate machine.

The actual data that the client sends is stored in the storage pools. Tivoli Storage Manager is unique because the storage pools can form a storage hierarchy made up of more than 500 supported devices. This allows for flexibility, longevity, and, most importantly, fast backups and fast restores.

Most organizations back up data initially to disk storage. This method enables hundreds of clients to back up at the same time. Then, based on policies, the data migrates in a fashion that expedites restores from tape or CD. When the data migrates, all data belonging to one client is moved together to the next storage pool in the hierarchy. By keeping all the data together, restores are faster because not as much tape positioning is required. This migration process can also accommodate movement to collocated tapes, which further expedites restores by storing only one user's data on them.

The environment can be firewall protected. However, clients can still use the GUI interfaces of a Tivoli Storage Manager client. Tivoli Storage Manager enables you to individually configure nearly every TCP port that it uses for communication. You can configure the following TCP ports:

- TCP/IP port

  To enable the backup/archive client, command-line administrative client, and the scheduler to run outside a firewall, the firewall administrator must open the port specified by the tcpport server option. This port is set on the client and the server using the tcpport option. The setting must be the same on the client and server. The default TCP/IP port is 1500.

- TCP/IP ports for the remote workstation

  The two TCP/IP ports for the remote workstation client must be open. Use the Web ports option in the remote workstations option file to specify these ports. If you do not specify the values for the Web ports option, the default zero (0) causes TCP/IP to randomly assign two free port numbers.

- HTTP port

  To allow the Web client to communicate with remote workstations across a firewall, the HTTP port for the remote workstations must be open. Use the httpport option in the remote workstation client options file to specify this port. The default client HTTP port is 1581.

To use the administrative Web interface for a server across a firewall, the Tivoli Storage Manager administrator must open the HTTP port for the server using the httpport option in the server options file. The default server HTTP port is 1580.

1.2.1 IBM Tivoli Storage Manager server

One of the principal architectural components of the Tivoli Storage Manager server is its built-in relational database. The Tivoli Storage Manager database was specially designed for the task of managing data, and it implements zero-touch administration.

All policy information, logging, authentication and security, media management, and object inventory are managed through this database.

Most of the fields are externalized through Tivoli Storage Manager high-level administration commands, SQL SELECT statements, or for reporting purposes, by using an ODBC driver. Obviously, this database is fully protected with software mirroring, roll-forward capability, and its own management and online backup and restore functions.

For storing the managed data, the Tivoli Storage Manager server manages a storage repository. The storage repository can be implemented in a hierarchy using any combination of supported media or magnetic or optical disk, tape, and robotic storage devices, which are locally connected to the server system or are accessible through a SAN. To take advantage of SAN technology, the Tivoli Storage Manager server has features that dynamically share SAN-connected, automated tape library systems among multiple Tivoli Storage Manager servers and provide (as an option) LAN or LAN-free and server-free backup.

1.2.2 IBM Tivoli Storage Manager backup/archive client

Data is sent to the Tivoli Storage Manager server using the IBM Tivoli Storage Manager backup/archive client and complementary Tivoli and non-IBM/Tivoli products. These products work with the Tivoli Storage Manager server base product to ensure that the data you store is managed as defined.

The Tivoli Storage Manager backup/archive client, which is included with the server, provides the operational backup and archival function. The client implements the patented progressive backup methodology, adaptive subfile backup technology, and unique record retention methods for backup and archive functions.

The backup/archive clients are implemented as multisession clients, which means that they are able to take advantage of the multithreading capabilities of modern OSs.

1.3 IBM Tivoli Storage Manager editions

This section discusses IBM Tivoli Storage Manager editions:

- 1.3.1, “IBM Tivoli Storage Manager Express” on page 9
- 1.3.2, “IBM Tivoli Storage Manager Basic Edition” on page 10
- 1.3.3, “IBM Tivoli Storage Manager Extended Edition” on page 12

1.3.1 IBM Tivoli Storage Manager Express

IBM Tivoli Storage Manager Express is aimed at two market segments, the small-to-medium business with a less sophisticated IT environment, or the enterprise department that does not require the full suite of Tivoli Storage Manager features.

Tivoli Storage Manager Express provides a subset of Tivoli Storage Manager features, focusing on backup and recovery for between 5 to 20 client machines. Tivoli Storage Manager Express offers the following features:

- Easy installation

  Tivoli Storage ManagerExpress takes less than one hour to install, configure, and start running backups.

- Simplified administration GUI

  A new GUI simplifies administration, and operational reporting is integrated. Client software deployment is also included.

- Fully upgradeable

  Tivoli Storage Manager Express can be upgraded to Tivoli Storage Manager Extended Edition.

- Disk-based incremental backup

  Client backups are achieved to disk storage pools on the Tivoli Storage Manager Express server. You have the option to use tape devices for longer term retention or off-site backups.

- Simplified tape management

  Use of traditional methods such as Grandfather/Father/Son backup sets simplify tape rotation, and all tape management is fully automated.

- Automatic configuration

  Clients are automatically configured with scheduled backups using industry best practices.

IBM Tivoli Storage Manager Express server is only supported in the Microsoft Windows environment.

Note: The Intel® Itanium® 64-bit processor is not supported.

For more information about the backup server and client requirements, refer to this site:

http://www-1.ibm.com/support/docview.wss?rs=3039&context=SSRQGY&uid=swg21230240

For more information about IBM Tivoli Storage Manager Express, refer to *Deployment Guide Series: IBM Tivoli Storage Manager Express*, SG24-7033, and visit the following Web site:

http://www-306.ibm.com/software/tivoli/products/storage-mgr-express/

1.3.2 IBM Tivoli Storage Manager Basic Edition

Tivoli Storage Manager Basic Edition contains a rich set of features and provides the core functions of backup, recovery, and archive management. Tivoli Storage Manager Basic Edition offers the following features:

- Progressive backup methodology

  Saves time and storage space by backing up only new and modified files. The progressive backup feature uses Tivoli Storage Manager's own relational database to track data wherever it is stored, delivering a direct one-step file restore. Progressive backup eliminates the requirement for traditional full-plus-incremental or full-plus-differential backup and restore procedures, commonly used by other storage management products.

- Tape resource sharing

  Enables multiple Tivoli Storage Manager servers to use the same tape library and drives, optimizing tape hardware asset utilization.

- Network-free rapid recovery

  Supports high-speed client data recovery directly from tape or optical devices. Recovery time is minimized by eliminating the use of network and central server resources.

- Dynamic multithreaded transfer

  Permits multiple clients to simultaneously transfer data to and from the same Tivoli Storage Manager server. Performance is boosted to more than three times the rate of a single-threaded session. The higher speed is achieved by transparently optimizing the number of data transfer sessions, based on available system resources.

- Adaptive differencing technology

  Changes the way data is backed up from the client. Using adaptive differencing, data is transferred to the server either by byte, block, or file level, based on the size of the file being backed up and the portion of the file that has changed since the last backup. Adaptive differencing technology supports all connectivity strategies, including LANs, WANs, SANs, Internet, and dial-up connections. Adaptive differencing was initially designed with mobile computer users in mind. However, other users who need to minimize data transmitted over the network can also benefit from the technology.

- Enterprise administration

  Simplifies centralized control across multiple Tivoli Storage Manager implementations without sacrificing network performance. Tivoli Storage Manager V5.4 is also powered by the Integrated Solutions Console (ISC), which provides a task-based GUI interface to Tivoli Storage Manager administrative tasks.

- Clustering

  Tivoli Storage Manager includes enhanced support for IBM High-Availability Cluster Multi-Processing (HACMP™), Microsoft Cluster Services (MSCS), Novell® Cluster Services (NCS), and Veritas Cluster Services (VCS) on Windows.

  Tivoli Storage Manager V5.4 has improved support for Small Computer System Interface (SCSI) and fibre-attached tape device failover on Microsoft Windows and UNIX®, and support for storage agents, library managers, and library clients as cluster members.

- LAN-free data transfer

  An optional module for Tivoli Storage Manager effectively exploits SAN environments by moving data transfers from the communication network to a SAN. Communication bandwidth availability is therefore improved, increasing service levels for users and clients.

- Hierarchical storage management

  An optional module for Tivoli Storage Manager automatically and transparently moves unused data files from online disk storage to offline tape storage. In the event that a file is accessed after it has been moved to offline storage, Tivoli Storage Manager transparently recalls the file.

- Library and device support

  Tivoli Storage Manager Basic Edition supports libraries with up to three tape drives and up to 40 cartridge capacity. Larger libraries can be accommodated but with only three devices and 40 slots enabled.

You can find more information about IBM Tivoli Storage Manager Basic Edition at this Web site:

http://www.ibm.com/software/tivoli/products/storage-mgr/

1.3.3 IBM Tivoli Storage Manager Extended Edition

The Extended Edition of IBM Tivoli Storage Manager expands on the features and possibilities of the Basic Edition described in the previous section.

Tivoli Storage Manager Extended Edition adds disaster recovery planning capability for the server, Network Data Management Protocol (NDMP) control for network-attached storage (NAS) filers, and support for larger capacity tape libraries and more tape drives.

You can find more information at:

http://www.ibm.com/software/tivoli/products/storage-mgr-extended/

Disaster Recovery Manager

The Disaster Recovery Manager (DRM) component of Tivoli Storage Manager Extended Edition provides disaster recovery for the Tivoli Storage Manager server and assists with disaster recovery for clients.

DRM offers various options to configure, control, and automatically generate a disaster recovery plan (DRP) file. The plan contains the information, scripts, and procedures required to automate restoration and help ensure quick recovery of data after a disaster. The scripts contain the commands necessary to rebuild the Tivoli Storage Manager server.

One of the key features of Tivoli Storage Manager and DRM is the ability to track media in all possible states, such as on site, in transit, or in a vault. The media movement features of DRM assist greatly with the daily tasks of sending disaster recovery media off site and receiving expired media on site for reuse. With these features you can quickly locate all available copies of data.

DRM functions help maintain business continuity through the following functions:

- Establishing and helping to automate a thorough *server* DRP: Clients can then subsequently restore their data from the server if required and can continue their daily backup procedures.
- Ensuring that vital site-specific information is available in the same plan.
- Automating vital recovery steps to return the Tivoli Storage Manager server and backup environment to normal operation.
- Managing and identifying off-site media required for recovery.
- Tracking and reporting destroyed systems in the event of a disaster.

- Storing client configuration information and assigning client recovery priorities.

With DRM you can recover at an alternate site, on a replacement system with a different hardware configuration, and with individuals unfamiliar with the applications. The DRP can be periodically tested to certify the recoverability of the server. The DRP can, and must, be recreated easily everyday so that it stays up to date. Figure 1-2 illustrates the main functions of DRM.

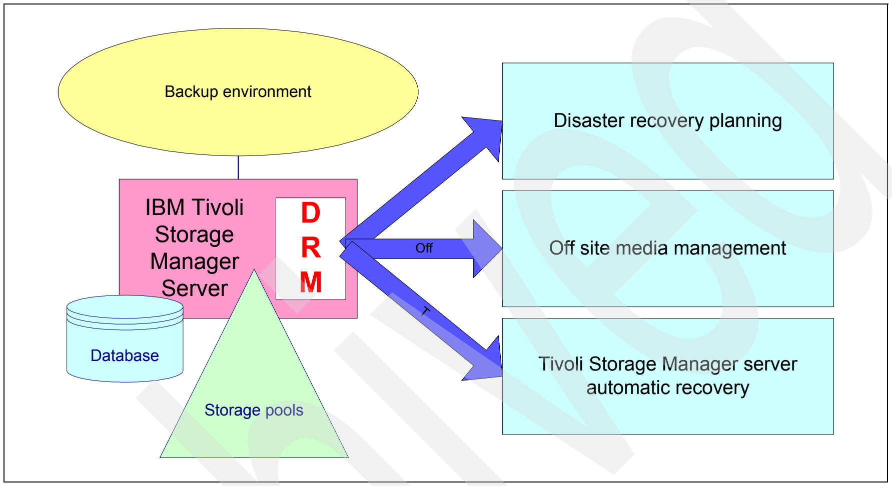

Figure 1-2 IBM Tivoli DRM functions

During a disaster, the following errors are commonly encountered:

- A DRP does not exist.
- The DRP was not tested, or if it was, it is now out of date.
- The testing team's skills were not sufficient to perform and evaluate testing.
- The disk volume definitions for the recovery site are not known.
- The location of recovery tapes is not known.
- It is not known which tapes are to be applied first.

DRM keeps track of all the vital information required to rebuild the Tivoli Storage Manager environment, including the following information:

- Current server configuration and its location
- Current Tivoli Storage Manager server database volumes (size, location, number)
- Recovery sequence
- Currency of the DRP
- Server and client machines configurations

- People to be contacted in the event of a disaster
- Location of the recovery media and the organization or persons responsible
- Point in time (PIT) to which the environment can be restored

During recovery from a disaster, DRM automates the following procedures to restore the Tivoli Storage Manager servers:

- Restore Tivoli Storage Manager server's key option files
- Copy files from alternate locations to production locations
- Initialize the Tivoli Storage Manager database and log volumes
- Match sizes and locations of the Tivoli Storage Manager database and log volumes
- Automatically launch restoration of the Tivoli Storage Manager database
- Track media required and availability
- Register installed Tivoli Storage Manager server features and return the server state to a valid license configuration
- Update Tivoli Storage Manager volume catalog information, regardless of whether volumes were destroyed during the disaster
- Rebuild Tivoli Storage Manager hierarchical storage configuration
- Restore destroyed volumes from those available where possible
- Recreate client backup environment

A detailed description, recovery scenario, and recovery plan built with DRM can be found in *Disaster Recovery Strategies with Tivoli Storage Management*, SG24-6844. The book also provides recommendations for and examples of using DRM to store client machine information in the DRM plan file for use during a client disaster recovery.

In summary, DRM systematically rebuilds the Tivoli Storage Manager server environment and ensures that current application data for the entire enterprise is available for recovery. You can do so automatically from a single scripted command.

NDMP support using IBM Tivoli Storage Manager

For NAS devices, Tivoli Storage Manager Extended Edition uses Network Data Management Protocol (NDMP) to perform high-performance, scalable backups and restores. NDMP-based backups and restores minimize network traffic and transfer data outboard of the Tivoli Storage Manager client and server. NDMP enables a full and differential file system image backup and restore of Network

Appliance™ file servers with OS Data ONTAP® V6.1.1 or higher and EMC Celerra systems running Dart V5.1.9.3 or higher.

NDMP support in Tivoli Storage Manager Version 5.5 includes the following hardware: IBM System Storage™ N series, Network Appliance, or EMC Celerra NAS file server. For supported models and operating systems, refer to:

http://www.ibm.com/tivoli/storage

Other NAS file servers support guidelines specified by Tivoli Storage Manager. NAS Vendors on the Ready for IBM Tivoli software list follow guidelines to implement NDMP as specified by Tivoli Storage Manager. File servers in the list have generally undergone tests to ensure they are compatible with Tivoli Storage Manager. For a list of NAS file servers that are certified through the Ready for IBM Tivoli software, visit:

http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html

For supported combinations of tape drives and tape libraries, refer to:

http://www.ibm.com/tivoli/storage

NDMP is supported one of the following operating systems:

- Windows 2003
- Windows 2003 64 bit SP1
- AIX
- Sun™ Solaris™
- HP-UX
- Linux® x86
- Linux pSeries®
- Linux zSeries®

Note: Tivoli Storage Manager supports NDMP on all the operating systems supported by the Tivoli Storage Manager Extended Edition Server except for the z/OS® operating system.

To perform NDMP backup and restore operations, Tivoli Storage Manager Extended Edition V5.4 or higher server is required.

You can use a Tivoli Storage Manager Extended Edition V5.4 or higher backup/archive client on Microsoft Windows, Sun Solaris (32 bit or 64 bit), or AIX (32 bit or 64 bit) to initiate volume-level backup and restore operations. To use a Web client to display files from a volume backup and select individual files to be restored, you must have a Windows Web client and a Web browser that can run on any platform.

The combination of file server model and operating system must be supported by the NAS file server. For more specifics, consult the product information for the NAS file server.

> **Note:** Tivoli Storage Manager supports NDMP Version 4 for all NDMP operations. Tivoli Storage Manager will continue to support all NDMP backup and restore operations with an NAS device running NDMP Version 3. The Tivoli Storage Manager server negotiates the highest protocol level (either Version 3 or Version 4) with the NDMP server when establishing an NDMP connection. If you experience any issues with Version 4, you might try using Version 3.

Extended library and drive support

Tivoli Storage Manager Extended Edition supports larger tape libraries, thus removing the 40-cartridge limit for library capacity and allowing more than four tape drives and 48 slots within a single library.

Database and application online protection

This feature protects a wide range of application data through the protection of underlying databases and application management systems that hold that data. This module automates data protection tasks and enables database and application servers to continue running their primary applications while they back up and restore data to and from offline storage. For more information, refer to:

http://www-306.ibm.com/software/tivoli/products/storage-mgr/product-links.html

Bare machine recovery

Cristie Bare Metal Restore (CBMR) combined with Tivoli Storage Manager functionality allows clients to recover a Microsoft Windows NT®, 2000, XP, or 2003 operating system to a new disk drive or RAID array or to a completely new machine using only a CD and a disaster recovery backup stored in the Tivoli Storage Manager server. This functionality is also supported for Linux, SUN Solaris, and HP-UX operating systems.

This feature backs up and automatically restores the OS structures required to rebuild the OS and data files. It schedules regular OS backups so that a recovery restores the latest information.

For more information, refer to:

http://www-306.ibm.com/software/tivoli/products/storage-mgr/cristie-bmr.html

1.4 Disaster recovery solutions

Disaster recovery must be part of every storage management solution. Tivoli Storage Manager Extended Edition offers serious disaster recovery management functions and solutions in addition to daily backup management.

In this section we discuss the types of contingency plans, disaster recovery plans (DRPs), and tier levels of disaster recovery solutions.

For more information about disaster recovery, refer to the *IBM System Storage Business Continuity: Part 1 Planning Guide*, SG24-6547.

1.4.1 Types of contingency plans

IT contingency planning represents a broad scope of activities designed to sustain and recover critical IT services after an emergency. IT contingency planning fits into a broader emergency preparedness environment that includes recovery planning and organizational and business process continuity.

In general, universally accepted definitions for contingency planning and related planning areas are not available. To provide a common basis for understanding IT contingency planning, this section identifies the following types of plans and describes their purposes:

- Business continuity planning

  Business continuity planning (BCP) confronts the likelihood of a disaster, how the disaster interrupts the business process, and how the business can continue in operation. The cause of the interruption does not matter; what matters is gaining management control and processing capacity immediately after the interruption.

- Business recovery plan

  The business recovery plan (BRP) (also business resumption plan) addresses the restoration of business processes after an emergency. The BRP is similar to the BCP, but unlike that plan, the BRP typically lacks procedures to ensure continuity of critical processes throughout an emergency or disruption.

- Continuity of operations plan

  The continuity of operations plan (COOP) focuses on restoring an organization's (usually the headquarters element) essential functions at an alternate site and performing these functions for up to 30 days before returning to normal operations.

- Incident response plan

  The incident response plan (IRP) establishes procedures to address cyber attacks against an organization's IT server and workstation systems. These procedures are designed to enable security personnel to identify, mitigate, and recover from malicious computer incidents, such as unauthorized access to a system or data, denial of service, or unauthorized changes to system hardware or software (such malicious logic as a virus, worm, or Trojan horse).

- Occupant emergency plan

  The occupant emergency plan provides the response procedures for occupants of a facility in the event of a situation posing a potential threat to the health and safety of personnel, the environment, or property. Such events include fire, hurricane, criminal attack, and medical emergency. Occupant emergency plans are developed at the facility level, specific to the geographic location and structural design of the building.

- Disaster recovery plan

  As suggested by its name, the DRP applies to major events that deny access to the normal facility for an extended period. Frequently, the DRP refers to an IT-focused plan designed to restore operability of the target system, application, data, or computer facility at an alternate site after an emergency. The plan's scope might overlap that of an IT contingency plan; however, the DRP is narrower in scope and does not address minor disruptions that do not require relocation.

1.4.2 Disaster recovery plan

The primary objective of disaster recovery is to protect the organization in the event that all or part of its operations and computer services are rendered unusable. Disaster recovery is the process of reacting to a disaster by being able to provide computing services from another location. In most cases, the countermeasures you employ to recover from a disaster are entirely different from the solution you use to achieve continuous availability.

A DRP establishes the procedures and actions to be performed when on the verge of a disaster. A disaster recovery solution can be the salvation of the enterprise and the origin of the whole disaster recovery initiative. For that reason, it is fundamental that you be able to manage an incident successfully, assisted with a high-quality plan.

To summarize, the DRP refers to a coordinated strategy involving plans, procedures, and technical measures that enable the recovery of IT systems, operations, and data after a disruption.

A DRP is a comprehensive statement of consistent actions to be taken before, during, and after a disaster. The planning must be documented and tested to ensure the continuity of operations and availability of critical resources in the event of a disaster.

The DRP applies to all (usually catastrophic) events that deny access to the normal facility for an extended period, including an orderly shut down of a facility to fix energy cables requiring an entire day. As stated previously, DRP refers to an IT-focused plan designed to restore operability of the target system, application, data, or computer facility at an alternate site after an emergency.

1.4.3 Seven tiers of recovery

The seven tiers of disaster recovery solutions (see Figure 1-3) offer a simple methodology to define your current service level, the current risk, and the target service level, and target environment.

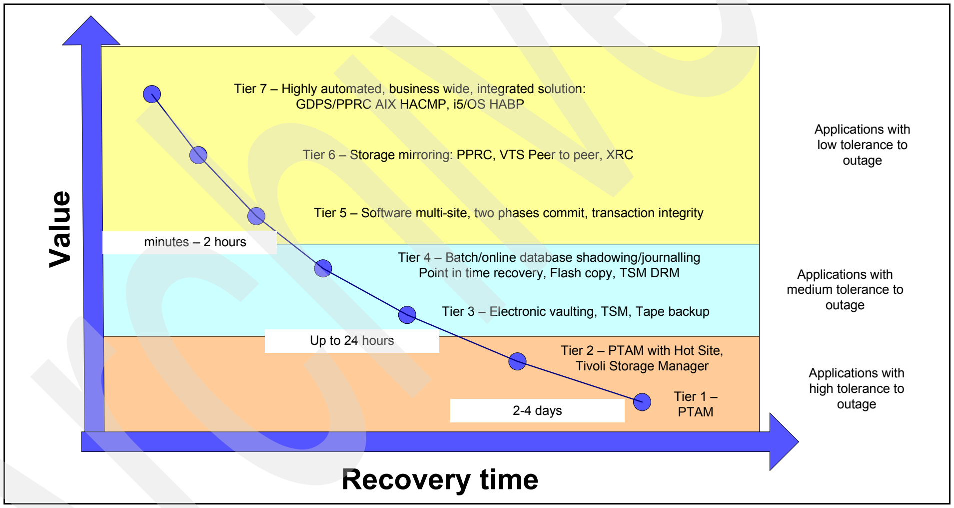

Figure 1-3 Disaster recovery tiers

- Tier 0: No off-site data

  Businesses with a Tier 0 disaster recovery solution have no DRP.

  – No saved information, no documentation, no backup hardware, and no contingency plan exist.

  – The length of recovery time in this instance is unpredictable. In fact, it might not be possible to recover at all.

- Tier 1: Data backup with no hot site

  Businesses that use Tier 1 disaster recovery solutions back up their data at an off-site facility. Depending on how often backups are made, these businesses are prepared to accept several days to weeks of data loss, but their backups are secure off site. However, this tier lacks the systems on which to restore data.

- Tier 2: Data backup with a hot site

  Businesses using Tier 2 disaster recovery solutions make regular backups on tape. This is combined with an off-site facility and infrastructure (known as a *hot site*) in which you can restore systems from those tapes in the event of a disaster. This tier of solution still results in the requirement to recreate several hours to days' worth of data, but it is less unpredictable in recovery time.

- Tier 3: Electronic vaulting

  Tier 3 solutions utilize components of Tier 2. Additionally, some mission-critical data is electronically vaulted. This electronically vaulted data is typically more current than data manually transported to a backup location. As a result less data is recreated or lost after a disaster occurs.

- Tier 4: Point-in-time copies

  Tier 4 solutions are used by businesses who require both greater data currency and faster recovery than users of lower tiers. Rather than relying largely on shipping tape, a solution commonly used on the lower tiers, Tier 4 solutions incorporate more disk-based solutions. Several hours of data loss is still possible, but it is easier to make such point-in-time copies with greater frequency than data replicated through tape-based solutions.

- Tier 5: Transaction integrity

  Tier 5 solutions are used by businesses with a requirement for consistency of data between production and recovery data centers. Little to no data is lost in such solutions; however, the presence of this functionality is entirely dependent on the application in use.

- Tier 6: Zero or little data loss

  Tier 6 disaster recovery solutions maintain the highest levels of data currency. They are used by businesses with little or no tolerance for data loss and who need to restore data to applications rapidly. These solutions have no dependence on the applications to provide data consistency.

- Tier 7: Highly automated, business-integrated solution

  Tier 7 solutions include all of the major components of a Tier 6 solution with the additional integration of automation. This enables a Tier 7 solution to ensure consistency of data above that which is granted by Tier 6 solutions. Additionally, recovery of the applications is automated, allowing for restoration of systems and applications much faster and more reliably than possible through manual disaster recovery procedures.

For detailed information, see the *IBM System Storage Business Continuity: Part 1 Planning Guide*, SG24-6547.

1.5 IBM Tivoli Storage Manager complementary products

Tivoli Storage Manager can be integrated with several optional applications that together form a powerful integrated storage management solution, including the following applications:

- “IBM Tivoli Storage Manager for Space Management” on page 22
- “IBM Tivoli Storage Manager for HSM for Windows” on page 23
- “IBM Tivoli Storage Manager for Storage Area Network” on page 23
- “IBM Tivoli Storage Manager for System Backup and Recovery” on page 24
- “IBM Tivoli Storage Manager for Databases product family” on page 24
- “IBM Tivoli Continuous Data Protection for Files” on page 27
- “IBM System Storage Archive Manager” on page 28
- “IBM Tivoli Storage Manager for Microsoft SharePoint” on page 29

For a full product listing, visit:

http://www.ibm.com/software/tivoli/products/

1.5.1 IBM Tivoli Storage Manager for Space Management

IBM Tivoli Storage Manager for Space Management provides hierarchical storage management (HSM) to automatically migrate rarely accessed files to alternative storage, without disrupting the most frequently used files in local storage. Migrated files are automatically and transparently recalled to primary storage when required by applications or users. Administrators and users are freed from manual file system maintenance tasks, and more online disk space is available for more important active data. Tivoli Storage Manager for Space Management can also help defer the requirement to purchase additional disk storage for clients, by making optimal use of available client storage.

Tivoli Storage Manager for Space Management offers increased scalability and performance through parallel migrations, improved candidate search, and optimized synchronization between the IBM Tivoli Storage Manager server and the HSM client.

IBM Tivoli Storage Manager for Space Management complements both IBM Tivoli Storage Manager and IBM Tivoli Storage Manager Extended Edition and is supported on AIX, HP/UX, Solaris, and Linux.

For more information, refer to:

http://www-306.ibm.com/software/tivoli/products/storage-mgr-space/

1.5.2 IBM Tivoli Storage Manager for HSM for Windows

IBM Tivoli Storage Manager for HSM for Windows is a new product that provides hierarchical storage management functionality to the Microsoft Windows platform. As with IBM Tivoli Storage Manager for Space Management, HSM for Windows automatically migrates rarely accessed files to alternative storage, without disrupting the most frequently used files in local Windows file systems. Similarly, migrated files are automatically and transparently recalled to their original location when required by applications or users.

Tivoli Storage Manager HSM for Windows allows various levels of granularity for migration of files. Files can be migrated individually, and file systems can be partially or fully migrated, based on a comprehensive set of policy options.

IBM Tivoli Storage Manager for Space Management complements both IBM Tivoli Storage Manager and IBM Tivoli Storage Manager Extended Edition.

For more information, refer to:

http://www-306.ibm.com/software/tivoli/products/storage-mgr-hsm/

1.5.3 IBM Tivoli Storage Manager for Storage Area Network

IBM Tivoli Storage Manager for Storage Area Network enables your SAN-connected Tivoli Storage Manager servers and client computers to make maximum use of their direct network connection to storage. This software extension enables both servers and client computers to make the bulk of their backup/restore and archive/retrieve data transfers over the SAN instead of the LAN, either directly to tape or to the Tivoli Storage Manager disk storage pool. This ability greatly reduces the impact of data protection on the LAN while also reducing CPU utilization on both client and server.

For computers running Microsoft Windows, some SAN configurations allow specific SAN devices to perform data movements directly to and from some tape devices, further reducing client and server CPU utilization.

Tivoli Storage Manager for Storage Area Networks complements and coexists with the standard library-sharing functionality of both Basic and Extended editions of the Tivoli Storage Manager server.

For more information see:

http://www-306.ibm.com/software/tivoli/products/storage-mgr-san/

1.5.4 IBM Tivoli Storage Manager for System Backup and Recovery

IBM Tivoli Storage Manager for System Backup and Recovery (SysBack™) provides a flexible backup method for AIX systems to help protect data and provide bare machine recovery capabilities. It offers a comprehensive system backup, restore, and reinstallation tool. SysBack is a simple-to-use, yet highly effective tool. Any feature can be executed from either the AIX command line or by using the SMIT menu interface.

For Microsoft Windows platforms, bare machine recovery can be achieved with the IBM Tivoli Storage Manager backup/archive client's Automated System Recovery capability.

In addition, Windows, Sun, and Linux bare machine recovery can be done with Cristie Bare Machine Recovery, which integrates directly with Tivoli Storage Manager to provide OS recovery for these platforms.

For more information see:

http://www-306.ibm.com/software/tivoli/products/storage-mgr-sysback/

1.5.5 IBM Tivoli Storage Manager for Databases product family

Tivoli Storage Manager provides data protection for a wide variety of applications, databases, mail, and hardware, ensuring that data is safe and secure no matter where it is located or how it is stored. These products interface directly with the applications using their backup-certified utilities and interfaces, simplifying online backup and restore procedures. We describe these products in the following sections.

IBM Tivoli Storage Manager for Application Servers

IBM Tivoli Storage Manager for Application Servers (formerly IBM Tivoli Data Protection for WebSphere® Application Servers) is a software module that works with Tivoli Storage Manager to better protect the infrastructure and application data and improve the availability of WebSphere Application Servers.

For more information see:

http://www-306.ibm.com/software/tivoli/products/storage-mgr-app-servers/

IBM Tivoli Storage Manager for Databases

IBM Tivoli Storage Manager for Databases is a software module designed to work with Tivoli Storage Manager to protect a wide range of application data through the protection of the underlying database management systems holding that data. IBM Tivoli Storage Manager for Databases exploits the various backup-certified utilities and interfaces provided for Oracle, Microsoft SQL Server, and IBM Informix®.

This same functionality is included in the IBM DB2 Universal Database™ package and Informix Dynamic Server, enabling them to work directly with Tivoli Storage Manager without the need for any additional modules.

For more information see:

http://www-306.ibm.com/software/tivoli/products/storage-mgr-db/

IBM Tivoli Storage Manager for Mail

IBM Tivoli Storage Manager for Mail is a software module for Tivoli Storage Manager that automates the data protection of e-mail servers running either Lotus Domino or Microsoft Exchange. This module utilizes the application programming interfaces (APIs) provided by e-mail application vendors to perform online backups and improve data-restore performance without shutting down the e-mail server. As a result, it can help protect the growing amount of new and changing data that must be securely backed up to help maintain Domino and Exchange application availability 24x7, 365 days a year.

For more information see:

http://www-306.ibm.com/software/tivoli/products/storage-mgr-mail/

IBM Tivoli Storage Manager for Enterprise Resource Planning

IBM Tivoli Storage Manager for Enterprise Resource Planning is a software module that works with Tivoli Storage Manager to better protect infrastructure and application data and improve the availability of SAP R/3® servers.

For more information see:

http://www-306.ibm.com/software/tivoli/products/storage-mgr-erp/

IBM Tivoli Storage Manager for Advanced Copy Services

IBM Tivoli Storage Manager for Advanced Copy Services (formerly known as IBM Tivoli Storage Manager for Hardware) is an optional software module for AIX that integrates with Tivoli Storage Manager Extended Edition. Tivoli Storage Manager for Advanced Copy Services protects mission-critical data that must be available 24x7, and integrates hardware and software-based snapshot capabilities with Tivoli Storage Manager and its Data Protection components for DB2 UDB, Oracle, and mySAP™.

Tivoli Storage Manager for Advanced Copy Services supports the following wide range of hardware:

- IBM Enterprise Storage Server® (ESS)
- IBM DS6000™
- IBM DS8000™
- SAN Volume Controller (SVC) and all IBM and non-IBM devices supported by the SVC. For a complete list, see:

  http://www.ibm.com/servers/storage/software/virtualization/svc/interop.html

Tivoli Storage Manager for Advanced Copy Services also provides the following functionality:

- IBM FlashCopy® support for ESS for Oracle
- FlashCopy support for ESS for DB2
- FlashCopy support for ESS for mySAP on DB2 UDB
- FlashCopy support for ESS for mySAP on Oracle
- Snapshot™ support for DS8000, DS6000, and SVC for DB2 UDB
- Snapshot support for DS8000, DS6000, and SVC for Oracle
- Snapshot support for DS8000, DS6000, and SVC for mySAP on DB2 UDB
- Snapshot support for DS8000, DS6000, and SVC for mySAP on Oracle
- Multiple snapshot versions managed by Tivoli Storage Manager policy
- Coordinated FlashCopy backup of multipartition DB2 UDB databases distributed across multiple host systems

Support of FlashCopy and snapshot functionality allows for "zero impact" backups and instant recovery. Data transfer to the Tivoli Storage Manager server is handled from a separate storage server, allowing the primary production data to remain online and undisturbed.

For additional information refer to:

http://www-306.ibm.com/software/tivoli/products/storage-mgr-advanced-copy-services/

IBM Tivoli Storage Manager for Copy Services

IBM Tivoli Storage Manager for Copy Services is a new optional module for Microsoft Windows that integrates with Tivoli Storage Manager or Tivoli Storage Manager Extended Edition. It is designed to leverage Microsoft Volume Snapshot Services (VSS) on Windows 2003. Tivoli Storage Manager for Copy Services provides functionality similar to Tivoli Storage Manager for Advanced Copy Services, but Tivoli Storage Manager for Copy Services supports Windows VSS for Microsoft SQL Server and Microsoft Exchange Server 2003 only.

Tivoli Storage Manager for Copy Services offers the following features:

- Single command-line interface (CLI) for performing legacy and VSS snapshot backup, restore, and query operations
- Single GUI for performing legacy and VSS snapshot backup, restore, and query operations
- Support for both hardware and software VSS providers that strictly adhere to Microsoft VSS provider requirements
- Support for a clustered Exchange environment

The Microsoft SQL VSS Integration module is used in conjunction with Data Protection for Microsoft SQL and the Tivoli Storage Manager Windows client. Together, they implement the Volume Shadow Copy Service (VSS) requestor interface to drive the Microsoft Volume Shadow Copy Service for backup and recovery of Microsoft SQL.

Full and copy backup types are supported, with granularity, at the Exchange Storage Group level. Tivoli Storage Manager policies manage backups, and they can be stored on the Tivoli Storage Manager server, local disks, or both. You can assign different policies for the different storage locations and backup types (full or copy).

As with Tivoli Storage Manager for Advanced Copy Services, zero impact backups and instant recovery enable the primary production data to remain online and undisturbed. Data movement to Tivoli Storage Manager storage can be offloaded to a secondary machine through a VSS hardware provider that supports transportable shadow copy volumes.

For more information see:

http://www-306.ibm.com/software/tivoli/products/storage-mgr-copy-services/

1.5.6 IBM Tivoli Continuous Data Protection for Files

According to industry surveys, almost 70% of corporate data exists on mobile computers or desktop machines, and less than 8% of it is backed up regularly. For mobile computer, desktop, and file server machines that contain important, critical, or sensitive data that is constantly being updated, a typical 24-hour backup cycle might not be sufficient to provide adequate data protection. The addition of Continuous Data Protection for Files provides a client machine with the capability to back up a file, transparently in real time, to a Tivoli Storage

Manager server as soon as the file is saved. The Tivoli Storage Manager server manages files backed up by this method in the same ways as other corporate data.

Tivoli Continuous Data Protection for Files was developed with mobile computer and desktop users in mind but can be applied to any client with a high rate of data change on its file systems.

Tivoli Continuous Data Protection for Files provides clients with true point-in-time recoverability. It is supported on AIX, Solaris, Linux, and Microsoft Windows platforms. For more information, see:

http://www.ibm.com/software/tivoli/products/continuous-data-protection/

1.5.7 IBM System Storage Archive Manager

IBM System Storage Archive Manager facilitates compliance with regulatory requirements. It helps manage and simplify retrieval of the ever-increasing amount of data that organizations must retain for strict records retention regulations. Many of the regulations demand the archiving of records, e-mails, design documents, and other data for many years, in addition to requiring that the data not be changed or deleted.

The IBM Tivoli Storage Manager existing policy-based data management capabilities help organizations meet many of the regulatory requirements of various government and industry agencies. Some new regulations require additional safeguards on data retention. IBM System Storage Archive Manager provides data retention policies that help meet these new regulations.

Data retention protection

IBM System Storage Archive Manager makes it extremely difficult to delete data before its scheduled expiration. Short of physical destruction to storage media or server, the deliberate corruption of data, or the deletion of the Archive Manager database, the System Storage Archive Manager does not allow data on storage managed by the Archive Manager server to be deleted before its scheduled expiration date. Content management and archive applications can apply business policy management for ultimate expiration of archived data at the appropriate time.

Features and functions

IBM System Storage Archive Manager hierarchical storage capabilities provides policies so that data is stored on the type of media that best meets that data's longevity, access speed, and cost needs.

Movement of the data from one media type to another (as media requires change or as new types of media become available) is achieved by *migration*. Migration automates moving the data to help ensure data longevity and also allows data to be stored on the type of media that best meets its speed of access and cost needs. System Storage Archive Manager includes the following features and functions:

- Expiration policies

  Expire the data when it is no longer required, thus freeing up the storage media, and providing cost effectiveness.

- Standard off-site data protection

  Off-site copies can be created on any of the hundreds of types of media supported, and like the primary copy, off-site copies are policy managed to allow for expiration.

- Archive client program

  Permits users to archive files from their workstations or file servers to archive retention-protected storage; users can also retrieve archived copies of files to their local workstations.

- Expiration and deletion suspension

  Enables you to place an unconditional hold on data, meaning data cannot be deleted or modified until you release the deletion hold.

- Event-based retention management

  Data is retained subject to a time interval that is calculated after a retention-initiating event occurs. The data then cannot be deleted until the time limit has expired. For example, you can specify records to be kept for a particular employee for one year after the employee leaves the organization.

- Data retention protection

  Data is not deleted until the retention criteria for the object is satisfied.

  For more information, visit the Web page:

  `http://www-306.ibm.com/software/tivoli/products/storage-mgr-data-reten/index.html`

1.5.8 IBM Tivoli Storage Manager for Microsoft SharePoint

IBM Tivoli Storage Manager for Microsoft SharePoint addresses the need for a fast, flexible, and real-time backup solution for the Microsoft Office SharePoint Server (MOSS) 2007, Microsoft OfficeShare PointPortal Server 2003, and Microsoft Windows SharePoint Services (V2.0-V3.0). The browser-based centralized interface provides full fidelity (including all metadata, securities, version histories, and customized layouts), which enables you to back up and restore at the item and subsite levels.

Integration with Tivoli Storage Manager meets client demand for a SharePoint-focused product. Tivoli Storage Manager for Microsoft SharePoint is aimed at reducing data loss and minimizing risks associated with organizations accessing, sharing and managing projects, and storing business-critical content and applications in SharePoint repositories.

Integration with Tivoli Storage Manager Extended Edition enables SharePoint users to perform hierarchal storage management, support multiple devices in their IT environment, use policy-based management for stored objects, and easily move data from disk to tape and then off site.

The capability of Tivoli Storage Manager for Microsoft SharePoint to support MOSS 2007 and to back up and recover at the item level sets it apart from any competing product in the market.

For more information see:

http://www-306.ibm.com/software/tivoli/products/storage-mgr-sharepoint/

1.6 Highlights of IBM Tivoli Storage Manager Version 5.5

In the following sections we describe the highlights of Tivoli Storage Manager Version 5.5. For full details, always refer to the announcement letter and to the installation and user guides for the relevant server. For full details, always refer to the announcement letter and to the *Installation and User Guides* for the relevant server. You can see the Version 5.5 announcement letter at:

http://www-01.ibm.com/common/ssi/rep_ca/6/877/ENUSZP07-0476/index.html

and for Version 5.4 at:

http://www-01.ibm.com/common/ssi/rep_ca/2/877/ENUSZP07-0102/index.html

The Tivoli Storage Manager documentation is available at:

http://publib.boulder.ibm.com/infocenter/tivihelp/

1.6.1 IBM Tivoli Storage Manager Version 5.5 new functions

Tivoli Storage ManagerVersion 5.5 includes the following new functions and capabilities:

- Enhanced disk utilization and performance
  - Sequential-access disk pool volume concurrent retrieve access
  - Sequential-access disk pool migration thresholds
- Enhanced support for Microsoft Windows
  - Supports online image backup on Microsoft Windows 64-bit operating systems
  - Open file support for Windows 64-bit operating systems
  - Supports files with up to 8184-character directory names
- Traditional backup and recovery enhancements
  - Additional device and operating system support for Tivoli Storage Manager for Advanced Copy Services Version 5.5
  - VSS Snapshot support for Microsoft SQL Server with Tivoli Storage Manager for Copy Services Version 5.5
- Server-managed encryption enhancements
  - Tivoli Storage Manager server-managed encryption keys for the backup/archive client. Tivoli Storage ManagerVersion 5.5 generates, encrypts, stores, and manages the encryption key in the Tivoli Storage Manager database.
  - Support for AIX Version 6.1 Encrypted File System (EFS) backup. Tivoli Storage Manager Version 5.5 backs up files in either clear text (decrypted by EFS) or in raw (encrypted) format.
- Additional server enhancements
  - Server enhancements include integration with VMware Consolidated Backup (VCB). The backup is performed from a VCB backup host, which can manage a virtual machine’s backup data as though it had been backed up by a Tivoli Storage Manager client running on the virtual machine.
  - Backup, archive, and space management support is provided for AIX workload partitions (WPAR). Tivoli Storage Manager Version 5.5 support enables you to back up and restore local partition file data within the global partition using the local partition namespace available within the global partition.
  - Server enhancements include restartable server-to-server export

– SAN device mapping for Virtual Tape Libraries (VTL) is included.
– Server enhancements include support for Internet Protocol Version 6 (IPV6).

The Tivoli Storage Manager Version 5.4 and Version 5.5 coexist, and both are enhanced with new functionality and device support. Thus the latest maintenance releases of Version 5.4 and the early releases of Version 5.5 concurrently introduce many changes. The following key enhancements are introduced in Tivoli Storage ManagerVersion 5.5:

- Enhanced security with Tivoli Storage Manager server-managed encryption keys for the client
- Efficient use of sequential-access disk pools
- Fault-tolerant export and import operations with restartable server-to-server export and import
- Fast, nondisruptive backups with Microsoft Exchange 2007 Volume Shadow Copy Service (VSS) support
- Integrated approach to protecting VMware with VMware Consolidated Backup (VCB) integration
- Broader support for application snapshots with:
  – Tivoli Storage Manager for Advanced Copy Services additional device and operating system support for DB2 snapshots
  – Tivoli Storage Manager for Copy Services VSS support for Microsoft SQL Server
- Support for backup, archive, and space management for AIX WPAR
- Support for Internet Protocol Version 6 (IPv6)
- Online image backup and open file support on Windows 64-bit operating systems
- AIX encrypted file system backup support

1.6.2 IBM Tivoli Storage Manager Version 5.5 client changes

The following key enhancements are available with the Tivoli Storage Manager backup/archive client V5.5:

- Transparent encryption support
- Secure socket layer (SSL) support
- Client-node proxy support
- TCP/IP Version 6 support

- AUDITLOGGING and AUDITLOGNAME options
- AIX JFS2 snapshot integration for snapshot-based image backup and snapshot-based file-level backup and archive
- AIX JFS2 extended attributes (EA) support
- Support for AIX Encrypted File System (EFS)
- AIX workload partition (WPAR) support
- Solaris ZFS™ support
- Linux Itanium 2 support
- Red Hat Enterprise Linux V5.0 support
- AIX V6.1 support
- IMAGEGAPSIZE option
- System services backed up as part of system state
- Dropped support for AIX V5.2
- Hierarchical storage management (HSM) for Microsoft Windows reconciliation and improved storage pool utilization

1.6.3 IBM Tivoli Storage Manager Version 5.5 technical changes

The following technical changes have been made to Tivoli Storage Manager V5.5:

- Calculating the migration thresholds for storage pools associated with sequential-access disk (file) devices uses a percentage of the storage pool's total data capacity.
- Concurrent access to volumes in storage pools with file device types is available.
- The following server processes are allowed shared read access to file volumes:
  - BACKUP DB
  - BACKUP STGPOOL
  - COPY ACTIVEDATA
  - EXPORT/IMPORT NODE
  - EXPORT/IMPORT SERVER
  - GENERATE BACKUPSET
  - RESTORE STGPOOL
  - RESTORE VOLUME

- The following commands are *not* allowed shared read access to file volumes:
  - AUDIT VOLUME
  - DELETE VOLUME
  - MIGRATION
  - MOVE DATA
  - MOVE NODEDATA
  - RECLAMATION
- Restartable server-to-server export is available.
- Tivoli Storage Manager V5.5 offers the capability to specify the TODATE and TOTIME parameters with EXPORT NODE or EXPORT SERVER commands.
- New support is available for Plasmon and UDO2 optical disk drives and media.
- Client and server authentication uses Secure Sockets Layer (SSL), with two new TCP/IP options SSLTCPPORT and SSLTCPADMINPORT.
- Tivoli Storage Manager V5.5 offers SAN discovery for nonroot users (AIX only).
- HP LTO-4 drive support and encryption is available.

  You must install the IBM RMSS Ultrium device drive to enable drive encryption with IBM LTO-4. IBM LTO-4 SCSI drives do not support encryption.
- 3592 drive encryption is enabled for HP-UX.
- Long file name support, up to 8704 bytes, is available.
- Tivoli Storage Manager V5.5 offers the capability to set a time zone on the z/OS server.
- z/OS server changes
  - The z/OS server no longer uses Server Virtual Machine for serialization. The z/OS server uses pthreads instead, which can improve performance.
  - The z/OS server is a POSIX-compliant UNIX System Services application.
- Tivoli Storage Manager V5.5 offers the capability to run SAN discovery for nonroot users.

For more information about the latest features and changes for Tivoli Storage Manager Version 5.5, refer to the *IBM Tivoli Storage Manager Versions 5.4 and 5.5 Technical Guide,* SG24-7447.

All Tivoli Storage Manager manuals can be found at:

http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp

1.7 IBM Tivoli Storage Manager V5.5 additional functionality overview

This section summarizes the status, new features, and changes for the following Tivoli Storage Manager components:

- "IBM Tivoli Storage Manager HSM for Windows Version 5.5" on page 35
- "IBM Tivoli Storage Manager for Space Management" on page 35
- "IBM Tivoli Storage Manager for Advanced Copy Services" on page 36
- "IBM Tivoli Storage Manager for System Backup and Recovery" on page 37

1.7.1 IBM Tivoli Storage Manager HSM for Windows Version 5.5

IBM Tivoli Storage Manager HSM for Windows provides space management for Microsoft Windows NTFS file systems. You can define file migration policies using the HSM for Windows GUI. File migration eligibility is determined by include and exclude policy criteria such as file type (extension) and various criteria related to the age of a file (creation, modification, last access). Tivoli Storage Manager HSM for Windows helps free administrators and users from file system pruning tasks. Tivoli Storage Manager HSM for Windows is designed to assist you to more effectively manage Windows NTFS disk storage by automatically migrating selected files based on policies you have established. The selected files are migrated to less expensive storage devices, while preserving Windows NTFS file accessibility.

See *IBM Tivoli Storage Manager HSM for Windows Administrator's Guide Version 5.5*, SC32-1773 and *Using the Tivoli Storage Manager HSM Client for Windows*, REDP-4126.

Tivoli Storage Manager HSM for Windows Version 5.5 offers the following enhancements:

- Added reconciliation function to reconcile migrated files with files that might have subsequently been moved or deleted from the source file server.
- Improved integration between Tivoli Storage ManagerHSM for Windows and the backup/archive client to help ensure that the current file content, even migrated files, is always stored in the backup.
- Enhanced performance.

1.7.2 IBM Tivoli Storage Manager for Space Management

The Tivoli Storage Manager for Space Management client for UNIX and Linux (the HSM client) migrates files from your local file system to distributed storage and can then recall the files either automatically or selectively. Migrating files to storage frees space for new data on your local file system and takes advantage of lower-cost storage resources that are available in your network environment.

Tivoli Storage Manager for Space Management is available for AIX JFS2 and GPFS™, Linux GPFS, Solaris VxFS, and HP-UX JFS file systems. Refer to the *IBM Tivoli Storage Manager for Space Management for UNIX: User Guide Version 5.5*, SC32-0148.

Tivoli Storage Manager for Space Management Version 5.5 includes the following enhancements:

- Increased performance and more efficient memory use
- Ability to handle more files in a single namespace
- General Parallel File System™ (GPFS) policy-driven migration
- Scout daemon dsmscoutd

  The scout daemon dsmscoutd has several new parameters to manage the daemon and to provide information about the file system.
- CFI database

  A Complete File Index (CFI) database is implemented to speed up candidate selection and reconciliation.
- Support for AIX V6.1 workload partition (WPAR)
- NFSv4 ACLs support for JFS2 and GPFS V3.2
- Support for HP-UX IA 64

1.7.3 IBM Tivoli Storage Manager for Advanced Copy Services

Data Protection for Snapshot Devices for DB2 Advanced Copy Services is a new component in IBM Tivoli Storage Manager for Advanced Copy Services. It provides the following features:

- Added support for the following operating systems when using N Series Snapshot devices:
  - Red Hat Enterprise Linux (RHEL) 4 Update 4 on System x™
  - Red Hat Enterprise Linux (RHEL) 5 on System x
  - SUSE® Linux Enterprise Server (SLES) 9 Service Pack 3 on System x
  - SUSE Linux Enterprise Server (SLES) 10 Service Pack 1 on System x
- Added support for N series with NAS attachment
- Added support for DB2 V9.5, which is the minimum supported level of DB2 for this component

- Ability to view data for snapshot backup operations through the Tivoli Storage Managerfor ERP Administration Assistant
- Added IPv6 support

For more information, refer to the release notes at:

http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp?topic=/com.ibm.itsmreadme.doc_5.5/relnote_acd550.html

1.7.4 IBM Tivoli Storage Manager for System Backup and Recovery

IBM Tivoli Storage Manager for System Backup and Recovery (also known as SysBack) provides system administrators and other system users with a simple, efficient way to back up and recover data from a command line or a SMIT menu-driven interface. SysBack enables you to recover all or part of the system. SysBack is also flexible; you can install one system installation image to another system with either identical or different hardware configurations, called *cloning*. You can use SysBack as a stand-alone product or in conjunction with a Tivoli Storage Manager server.

Refer to the *IBM Tivoli Storage Manager for System Backup and Recovery 6.1 Installation and User's Guide*, SC23-6543.

Tivoli Storage Managerfor System Backup and Recovery offers the following enhancements in Version 6.1:

- Supports the system boot process with CD and DVD devices when recovering the system using data stored in a Tivoli Storage Manager server.
- Supports the use of the Tivoli Storage Manager backup/archive client file-by-file backup data to recover the system during bare machine recovery processing.
- An LVM-only backup option is available for use together with Tivoli Storage Manager backup/archive client backups during system install processing.
- LVM-only backups that back up just the LVM information for the complete system to CD and DVD, tape, virtual device, Tivoli Storage Manager server, and directory devices. These backups can be used in the following ways:
  - As boot media
  - To recreate LVM structures in normal mode
  - As part of the system installation and recovery process when restoring the system using backups generated by the Tivoli Storage Manager backup/archive client.
- Provides JFS2 snapshot backups that create snapshot copies of JFS2 file systems and then use SysBack backup commands to back up those copies.

- Offers new options to limit the amount of volume group, logical volume, file system, and physical disk information collected during a system-level backup.
- Offers a new backup option to allow backups to continue even when a missing or an invalid backup object is specified to the backup command.
- Supports enhanced network install debugging options.
- Provides automatic backup and restore process logging options.
- Provides an activity logging option for an additional level of process information.
- Includes a SMIT menu panel to collect detailed environment information for your system and the SysBack product.

1.8 Summary

In this chapter, we provided an overview of IBM Tivoli Storage ManagerV5.5 and other complementary storage management products from IBM Tivoli. We described the architecture of Tivoli Storage Managerand disaster recovery concepts that are vital for understanding storage management. We also listed new functions available with Tivoli Storage Manager Version 5.5. We describe more of these functions in subsequent chapters of this book.

2

IT environment

This chapter discusses the architectural and sizing considerations that you must address before deploying IBM Tivoli Storage Manager. In addition, we provide solutions to specific architectural scenarios.

In this chapter, we consider the following topics:

2.1 IBM Tivoli Storage Manager system requirements

In this section, we provide an overview of the hardware and software requirements for installing and configuring Tivoli Storage Manager V5.5. For detailed information see the following Web site:

http://www-306.ibm.com/software/tivoli/products/storage-mgr/platforms.html

The discussion includes:

- 2.1.1, “IBM Tivoli Storage Manager supported server OS and hardware” on page 40
- 2.1.2, “IBM Tivoli Storage Manager supported clients OS and hardware” on page 41
- 2.1.3, “IBM Tivoli Storage Manager API client” on page 41
- 2.1.4, “IBM Tivoli Storage Manager supported storage devices” on page 41
- 2.1.5, “IBM Tivoli Storage Manager Administration Center requirements” on page 42
- 2.1.6, “Server compatibility with storage agents and library clients” on page 42
- 2.1.7, “IBM Tivoli Storage Manager platform support update” on page 43

2.1.1 IBM Tivoli Storage Manager supported server OS and hardware

The Tivoli Storage Manager server V5.5 is supported on the following operating systems:

- HP-UX PA
- Microsoft Windows
- Sun Solaris
- OS/390® z/OS
- Linux x86_64
- Linux x86
- Linux zSeriesV5
- Linux on Power

2.1.2 IBM Tivoli Storage Manager supported clients OS and hardware

The Tivoli Storage Manager backup/archive client V5.5 is supported on the following operating systems:

- AIX
- HP/UX
- HP/UX Itanium
- Linux for IA64
- Linux on Power - iSeries/pSeries
- Linux on x86 / x86_64
- Linux zSeries
- Macintosh
- Novell NetWare
- OS/390, zSeries UNIX System Services
- OS/400®
- Sun Solaris
- Solaris x86 / x86_64
- Microsoft Windows
- Citrix Presentation Server 3.0 for Windows 2000 and 2003

2.1.3 IBM Tivoli Storage Manager API client

This section discusses all supported Tivoli Storage Manager clients, except Macintosh.

2.1.4 IBM Tivoli Storage Manager supported storage devices

The following storage devices for different platforms are supported by Tivoli Storage Manager:

- AIX, HP, SUN, and Windows

  http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_AIXHPSUNWIN.html
- iSeries®

  http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_iSeries.html
- Linux

  http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_Linux.html

2.1.5 IBM Tivoli Storage Manager Administration Center requirements

For detailed information about system requirements for the Integrated Solution Console (ISC) and the Administration Center, see:

http://www-1.ibm.com/support/docview.wss?uid=swg21286856

2.1.6 Server compatibility with storage agents and library clients

The following interoperability guidelines apply to the LAN-free storage agent, library client, and server for IBM Tivoli Storage Manager:

- A given server can never support an “up-level” storage agent.” A storage agent is considered to be up-level if it is at a higher version, release, or fix level. For example, a V5.5.0 server does not support a V5.5.1 storage agent. This point is illustrated in Table 2-1.
- A problem might occur between a Tivoli Storage Manager server and a storage agent or library client that is not at the latest modification level for that version and release of the server. If a problem occurs, you might have to upgrade the storage agent or library client to the server’s modification level.
- A new function that is available in a specific version, release, and modification level of the server is not supported by an earlier version, release, and modification level of the storage agent. Table 2-1 shows the storage agent and library client compatibility with Tivoli Storage Manager servers configured as library managers.

Table 2-1 Compatibility matrix

Tivoli Storage Manager server (library manager) version	Tivoli Storage Manager library client and storage agent versions
V5.5	V5.5, V5.4, and V5.3
V5.4	V5.4 and V5.3

For more information about storage agent and library client compatibility with Tivoli Storage Manager servers V4 and V5, see the following Web site:

http://www-1.ibm.com/support/docview.wss?uid=swg21302789

If a problem arises with the LAN-free support or library client in a given client environment, you must first consider whether or not the storage agents can be upgraded to match the version and release of the server.

2.1.7 IBM Tivoli Storage Manager platform support update

Select support for three operating systems was added to Tivoli Storage Manager V5.4 on April 30, 2008, via special V5.3.6 backup/archive clients and V5.3.6.3 storage agents. These clients and storage agents are supported in Tivoli Storage Manager V5.4 for use with Tivoli Storage Manager V5.4 and V5.5 servers:

- Windows 2000
- Solaris 8
- Linux x86 RHEL 3

Tivoli Storage Manager V5.3.6 client download page:

http://www.ibm.com/support/docview.wss?uid=swg24019078

TSM client V5.3.6.3 storage agents download page:

ftp://ftp.software.ibm.com/storage/tivoli-storage-management/maintenance/storage-agent/v5r4/special_platform_support/

Support ends 90 days after vendor-extended defect support ends for the operating system, as listed here:

- Windows 2000: July 2010
- Solaris 8: March 2009
- Linux x86 RHEL 3: October 2010

To align with the current Hewlett-Packard support plan, Tivoli Storage Manager added support for HP-UX 11i V1 to Version 5.4 in March 2008, using the V5.4.2 backup/archive client and the V5.4.3 server and storage agent.

2.2 Solution architecture considerations

The planning phase of a IBM Tivoli Storage Manager implementation is very important. The chosen design and architecture must support the needs of your business. Design decisions have a direct impact on both the functionality and performance, as well as the cost of the solution. The time you spend considering the various requirements and features of Tivoli Storage Manager is valuable and ensures that the right and most cost-effective solution is reached. This planning phase is especially important for large, complex implementations because seemingly insignificant choices can greatly affect cost and performance due to the sheer number of Tivoli Storage Manager clients (licenses), amount of storage space needed (disk or tape), and the choice of transport layer (LAN, WAN, or SAN). The cost of licenses and hardware is easy to take into account. The real challenge is to identify hidden expenses, such as power and cooling costs or the administration costs of the fully implemented solution.

Figure 2-1 illustrates the relationship between increased performance or complexity of a Tivoli Storage Manager solution and expenses for hardware, software, and administration.

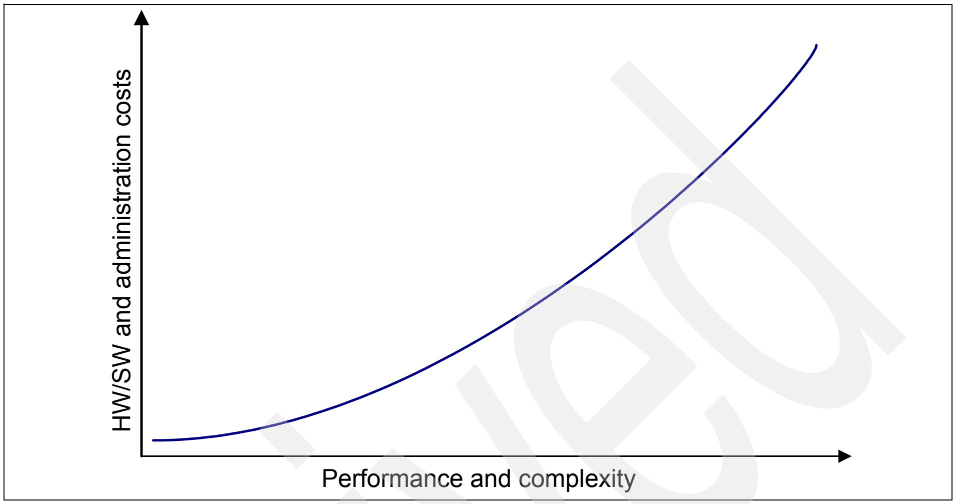

Figure 2-1 Relationship between performance, complexity, and costs

2.2.1 Backup and restore service level agreements

You first need to obtain the information about service level agreements (SLAs) describing requirements for the backup, restore, and retention time of your backup data, before you begin to plan your IBM Tivoli Storage Manager environment. This information dictates the design and therefore the features needed to support the SLAs.

Backup solutions must always be designed for restore, not the other way around.

During the design process, it is critical that you as the Tivoli Storage Manager administrator obtain input and business requirements for backup and restore time frames. You use this information in the design process to determine the type of equipment required to support the solution.

The following list describes the data that, ideally, you must obtain for each machine that has Tivoli Storage Manager client installed.

- Client name: Useful for documentation process
- Client OS: What types of OSs are necessary to support in the environment and what features are available on the client?
- Network speed: How fast can the system transmit data over the network?

- Backup window: How much time do you have to back up the system on a nightly basis?
- Restore window: How long do you have to restore the system? This consideration includes restoring the whole system as well as just a few files.
- Estimated growth: How fast is the data on the system growing?
- Amount of data: How much data is on the system?
- Type of data: What type of data is on the system? Is it file system data or database data? How many files?
- Retention policy: How long must data from the system be kept? Is both a backup and an archive needed?
- Nightly change: What percentage of the data changes on a nightly basis? This issue concerns systems that mostly use progressive backups.

If the Tivoli Storage Manager environment is small enough, this information is fairly easy to gather. In a large, dynamic environment, however, the information is challenging to obtain. You can facilitate the process of gathering of this information in several ways.

First, note that several of the inputs lend themselves to be applied globally throughout the Tivoli Storage Manager environment. SLAs and retention policies can be established and promoted on a company-wide basis. In addition, you can estimate nightly changes based on the backup methodology used. For example, it is safe to assume that most incremental file system backups include only about 10% of the total data on a system. And, of course, full backups of databases and file systems include 100% of the total data.

Another option is to use one of several tools on the market that can help facilitate gathering server-specific information about storage usage and growth. IBM TotalStorage Productivity Center for Data (TPC for Data) is one such product. TPC for Data scans computers and provides a centralized manager that you can use to track individual server file and disk usage, growth patterns, and total data usage throughout an environment. You can use this information as input to the architectural process.

You can locate up-to-date information about IBM TotalStorage Productivity Center for Data at:

`http://www-03.ibm.com/systems/storage/software/center/data/`

As you go through the architectural process, however, you likely find it impossible to gather all required information for a Tivoli Storage Manager architecture with 100% accuracy. If you cannot obtain an input or it is impractical to get it, substitute it with a reasonable assumption. For example, suppose you do not know how fast data is growing in your environment. A typical assumption is 10%-30% per year.

It is important to revisit both your concrete inputs and your assumptions on a regular basis. Doing so enables you to re-evaluate and adjust your assumptions, refresh your real data, and determine how the Tivoli Storage Manager architecture must be updated to account for these changes. Just as inputs that go into producing an architecture are rarely static, the architectural solution you produce based on these inputs require ongoing revision and refinement.

2.2.2 IBM Tivoli Storage Manager server hardware

The Tivoli Storage Manager server hardware is critical to the overall Tivoli Storage Manager infrastructure and design. Even though Tivoli Storage Manager is a piece of software, its architecture is intimately tied to the hardware on which it runs. In its simplest terms, a Tivoli Storage Manager server is used for running the database that stores the metadata, moving data from clients to the storage subsystem, and moving data from a storage subsystem to another storage subsystem. The speed of the CPU, disk subsystem, backplane, and other components of the Tivoli Storage Manager server are critical factors in determining how much data can be moved in a given time period. It is important to understand how specific hardware influences the solution and where it makes little or no difference what you choose.

Building Tivoli Storage Manager servers with numerous CPUs, disk, or RAM resources does not always guarantee a system that meets performance expectations. When choosing a Tivoli Storage Manager server platform and model, the most important aspect is the overall data workload the system has to support, meaning I/O, and determining how much data moves in and out of the system and if the system is able to support it.

Best practices:

- Choose a Tivoli Storage Manager server platform and model based on cost, expendability, current OS knowledge, and the ability to move data.
- I/O demanding Tivoli Storage Manager solutions tend to run better on UNIX systems, and less-demanding solutions are easier to administer on Microsoft Windows systems.
- Ensure that the bus and backplane of the system you choose is capable of handling the data workload and the number of expansion cards necessary to connect I/O devices.
- Always make sure you have enough expansion slots to separate disk I/O from tape drive I/O. Running both types on the same Fibre Channel cards can produce unpredictable results.
- Consider higher CPU counts and speeds for systems that make extensive use of Gigabit Ethernet. CPU is less important for systems moving most of the data across a SAN through Fibre Channel connections.
- Try to obtain performance data from a Tivoli Storage Manager system similar to the system you are planning to implement. Real-world experience is invaluable.

When it comes to the specific task of running Tivoli Storage Manager server processes, one OS seldom has inherent advantages over another. All supported platforms can do an adequate job running the Tivoli Storage Manager server code as long as the hardware is sized appropriately. In addition connecting particular Tivoli Storage Manager clients to a similar server (for example, a Microsoft Windows client to a Windows server) offers no inherent advantage; Tivoli Storage Manager is designed to connect any platform client to any platform server. Generally, 32-bit servers deliver as much performance as 64-bit systems. The main advantage of 64-bit systems is the ability to address more than 4 GB of RAM.

An overall discussion of server hardware platforms for the myriad systems that Tivoli Storage Manager supports (AIX, Solaris, HP-UX, Microsoft Windows, IBM z/OS, and IBM OS/400) is beyond the scope of this book. However, you must keep in mind the following general rules as you choose distributed server hardware:

- A Tivoli Storage Manager server usually has two to four CPUs per Tivoli Storage Manager instance. Processing the TCP/IP stack in a LAN backup environment takes significant CPU, so the more high-speed network cards (for example, Gigabit Ethernet) you have, the more CPU you require.
- Tivoli Storage Manager does not use an excessive amount of RAM. A good starting point for any modern distributed server OS is 2 GB of RAM. Add 1 to 2 GB for each Tivoli Storage Manager instance that the system hosts.
- Backplane speed and number of expansion slots are critical considerations. Remember that most of what a Tivoli Storage Manager server does is take in data through the network or SAN and write it to storage devices such as disk and tape. Use your data workload calculations to determine how many LAN and SAN adapters you require and how many tape drives you have to support. Ensure that the system you choose has enough slots to meet the requirement for LAN, SCSI, and Fibre Channel cards. Make sure you are not shortchanging yourself; “cheap” systems and “expensive” ones can differ widely in bus speed, expandability, and I/O capabilities.
- It is important to choose the right hardware. Tivoli Storage Manager does not directly support changing hardware platforms (for example, changing from x86 to POWER6™). The database structure is different on the various hardware platforms and therefore cannot be ported between platforms. The only way to change a platform is to export all stored data across a network to another Tivoli Storage Manager server. This process can, however, be tedious and costly.

2.2.3 IBM Tivoli Storage Manager server instances and placement

As part of the design process, you must decide how many instances of the Tivoli Storage Manager server are to be used and the placement of each. An instance of a Tivoli Storage Manager server is defined as the database and its corresponding recovery log. Tivoli Storage Manager server instances can share the same computer hardware and libraries using server-to-server communications and library sharing, or they can be installed on completely separate hardware.

Best practices:

- Choose the number of and location for Tivoli Storage Manager server instances based on workload, network infrastructure, database size, and the ability of the hardware to connect to storage.
- If multiple instances are required due to database size, consider running multiple instances of Tivoli Storage Manager on the same physical server.
- Consider using library sharing if you are running multiple Tivoli Storage Manager instances. Library sharing greatly reduces library and media administration and helps simplifies the solution.

You must carefully plan and monitor the sizes of the Tivoli Storage Manager database and recovery log to ensure that the system operates efficiently. You can estimate the size of the database using certain formulas to keep client metadata at a reasonable size so that the system does not bog down during backups and other database-related transactions, such as extracting information from the database for reporting purposes using SQL SELECT commands.

Database size and performance are the primary drivers of the number of server instances a Tivoli Storage Manager solution requires. If the size of the Tivoli Storage Manager database is such that performance expectations cannot be met, multiple instances of Tivoli Storage Manager must be used to accommodate the performance requirements.

The location of Tivoli Storage Manager servers in the enterprise is primarily determined by the network connection speed between sites. For simplicity, create as few Tivoli Storage Manager servers as possible. However, sometimes a single server or a few central Tivoli Storage Manager servers are not enough to support the backup and restore requirements and the SLAs when considering the whole network infrastructure, client sites, and Tivoli Storage Manager client placements. You must consider workload calculations and WAN transport speeds to determine whether backups and restores are feasible across a WAN connection. If not, then local Tivoli Storage Manager servers have to be deployed so that the data does not have to flow across WAN connections, which are too slow to support backing up to the central Tivoli Storage Manager server.

When determining the feasibility of Tivoli Storage Manager clients accessing the Tivoli Storage Manager server over a WAN, consider both backup and restore workloads. For file system backups, Tivoli Storage Manager can do incremental forever backups, which significantly reduce the amount of data that has to be moved. Even though incremental forever backups can make day-to-day WAN backup feasible, full system restores might not be possible to perform within an adequate time frame.

2.2.4 IBM Tivoli Storage Manager data transport possibilities

Tivoli Storage Manager is a network backup and archive application. As such, data moves from the clients to the server (or directly to storage) over some sort of network transport. Several popular choices for transports are available, including LAN and SAN. Although you might not be able to provide detailed networking infrastructure recommendations as a Tivoli Storage Manager administrator, you need to evaluate the different types of transports and make recommendations on the most appropriate transport (LAN versus SAN) and on the network infrastructure that is needed to meet backup and restore requirements.

LAN and WAN backups

The performance of your network backup solution is no better than the performance of your network. You must consider both theoretical and real performance of your network when you are designing a Tivoli Storage Manager solution.

Best practices:

- Assume 40%-80% of total theoretical throughput for a TCP/IP adapter or protocol.
- Be careful of the CPU usage and real throughput of faster TCP/IP protocols (such as Gigabit Ethernet). They often do not perform close to the 80% rule.
- Measure LAN throughput outside of Tivoli Storage Manager with simple protocols such as FTP. This measurement provides a true picture of how a network is performing.

Figure 2-2 shows a diagram of the flow of data during LAN and WAN backups. Notice that the data flow and metadata flow are the same. Metadata is the information about what is being backed up and not the actual files.

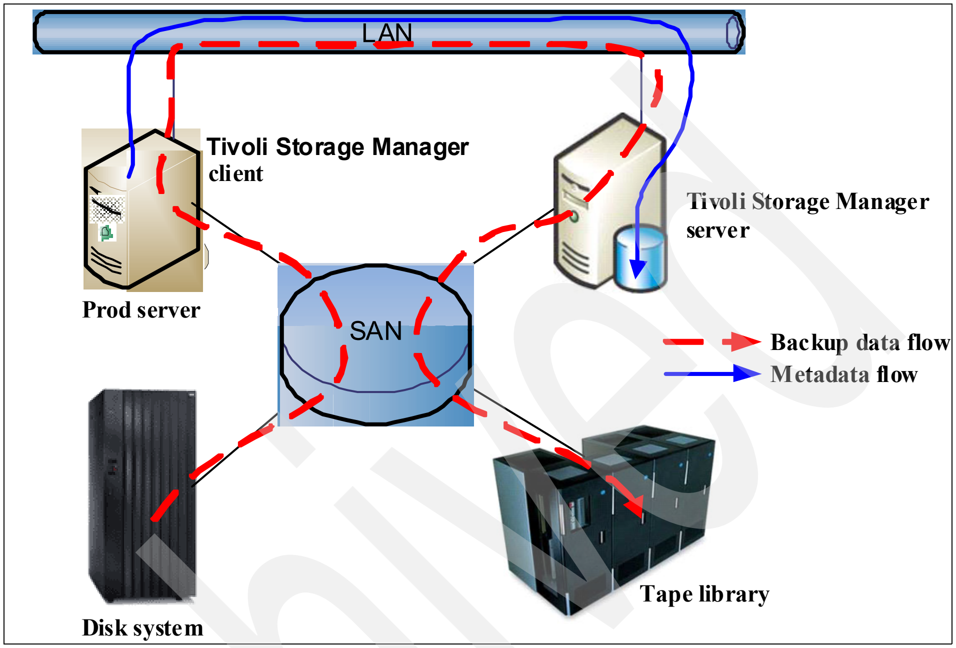

Figure 2-2 LAN and WAN backup

You can calculate theoretical network throughputs. For example, using Fast Ethernet and a 40% efficiency, the following calculation of the theoretical network throughput applies:

$$\frac{100Mb}{s} \times \frac{1B}{8b} \times \frac{1GB}{1024MB} \times \frac{60s}{1\min} \times \frac{60\min}{1Hr} \times .4 \approx 18\frac{GB}{Hr}$$

Using this type of calculation, you can calculate the anticipated throughput for any network transport. Table 2-2 shows the most common throughputs.

Table 2-2 Common network throughputs

Technology	Mbps	Assumed speed
T1	1.54	0.5 GB/hr @ 80% efficiency
Ethernet	10	2 GB/hr @ 40% efficiency
T3	45	16 GB/hr @ 80% efficiency

Technology	Mbps	Assumed speed
Fast Ethernet	100	18 GB/hr @ 40% efficiency
Gigabit Ethernet	1,000	180 GB/hr @ 40% efficiency

Using Table 2-2 on page 51, your own calculations, and real-world testing, you can validate the amount of data you need to move, back up, and restore (that is, to the backup and restore workload) against your SLAs and your network transports. Obviously, if the network performance does not meet your SLAs, either the SLAs must be relaxed or the network must be improved.

LAN-free backup or SAN backup and restore

IBM Tivoli Storage Manager for Storage Area Networks enables Tivoli Storage Manager backups and restores to be sent over the SAN directly to tape drives instead of being sent over the LAN to the Tivoli Storage Manager server. In Tivoli Storage Manager, this process is called *LAN-free*. However, “LAN-free” is somewhat of a misnomer because only the data moves across the SAN; the metadata (information about what is being backed up) still moves across the LAN to the Tivoli Storage Manager server. The LAN-free setup relieves the LAN, Tivoli Storage Manager server, and the client processor. Instead, the load moves to the SAN and the Fibre Channel cards in the client (which usually uses less CPU than the processing of the TCP/IP stack in the LAN adapters).

Using LAN-free backup does not always ensure fast backup and restore. The data flow moves from the LAN to the SAN, so the LAN is no longer the bottleneck. However, the bottleneck is just moved (not removed), and often the new bottleneck is the disk subsystem that holds the client data. This disk subsystem must be able to deliver adequate performance to feed the tape drives directly.

As seen in Table 2-3, SAN throughput greatly outperforms LAN throughput.

Table 2-3 Common SAN throughputs

Technology	Mbps	Assumed speed
Fibre Channel 2 Gbit	2,000	700 GB/hr @ 80% efficiency
Fibre Channel 4 Gbit	4,000	1,400 GB/hr @ 80% efficiency

Best practices:

- Transporting backup data across a SAN is usually 2-5 times faster compared to transporting backups across a LAN.
- When planning for LAN-free, remember that file metadata must still be sent over the LAN.
- Tivoli Storage Manager clients with lots of small files, for example, file server type systems, often perform better over a LAN than over a SAN. LAN-free is most appropriate and cost effective for systems with large files and large amounts of data (such as database and mail servers).
- Evaluate the cost of Gigabit Ethernet. It might be more cost effective and easier to manage than a backup SAN.
- Administration and license costs are higher when using LAN-free backup.

Figure 2-3 shows a typical LAN-free setup and data flow.

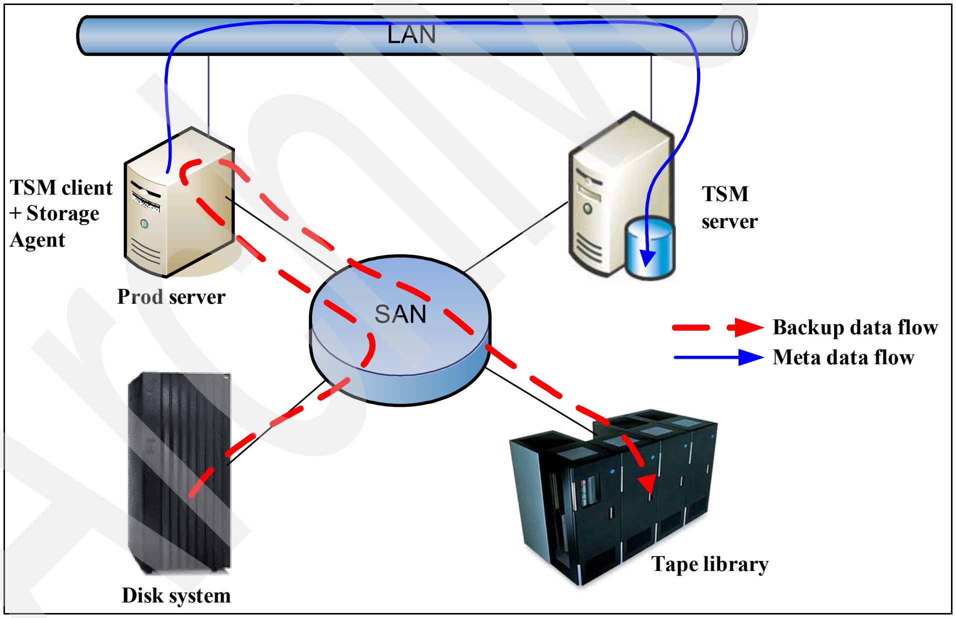

Figure 2-3 LAN-free backup

When designing a Tivoli Storage Manager solution, it is important to understand that the metadata can add overhead to a LAN-free backup. Even though the data is streaming over the SAN to the tape or disk, the metadata still goes to the Tivoli Storage Manager server using any LAN protocol and speed that are available. This bottleneck is most apparent when backing up millions of small files.

Experience has shown that systems with large files (greater than 10 MB on average) and large amounts of data (greater than 50 GB) are the best candidates for LAN-free backup, while systems with numerous small files (average size in KB) often perform much better using traditional LAN-based backups. Examples of the former include database and mail systems and enterprise resource planning (ERP) systems; examples of the latter are traditional file- and print-type servers.

Split-mirror backup

An alternative setup of the previously mentioned LAN-free backup includes a second client that mounts a mirror of the production client's data and sends it across the SAN directly to the tape drives. The mirror is generated by OS mirroring or a hardware-assisted instant copy function (found in many of today's high-end storage systems). This setup is especially useful when performance of the production system is crucial and no performance impact is allowed during backup. The functionality is the same as for the LAN-free setup, except for the second client, which, in this example, is in control of sending the data (see Figure 2-4 on page 55). The setup enables Tivoli Storage Manager to perform a backup without the data stream leaving the SAN or passing through the production system.

Best practices:

- Use the split-mirror backup when "zero impact" is essential for your production system.
- Administration and license costs of a split-mirror backup are high.

Be aware that split-mirror backup functionality is much harder to implement and maintain in daily production than LAN-free backup. This difficulty is caused by the complexity of the scripts that break off and reattach the data mirror. It can, however, be automated using the Tivoli Storage Manager for Advanced Copy Services complementary product by coupling the N Series Snapshot or FlashCopy for IBM Enterprise Storage Server with Tivoli Storage Manager and its database protection capabilities for DB2, Oracle, and SAP R/3 databases.

Using FlashCopy and Snapshot functionality enables "zero impact" backups and instant recovery. You create a shadow of your original data volume on your disk subsystem (for instance, IBM DS8000 or N Series, respectively). You can create this volume, or at least the pointers, in just a few seconds. Your application does not see any big impact during this creation. This just-created flash volume can be also used for a fast restore because it needs only to be "flashed back." This flash volume is used for copying the data to tape. You are most likely to copy the data LAN-free to tape.

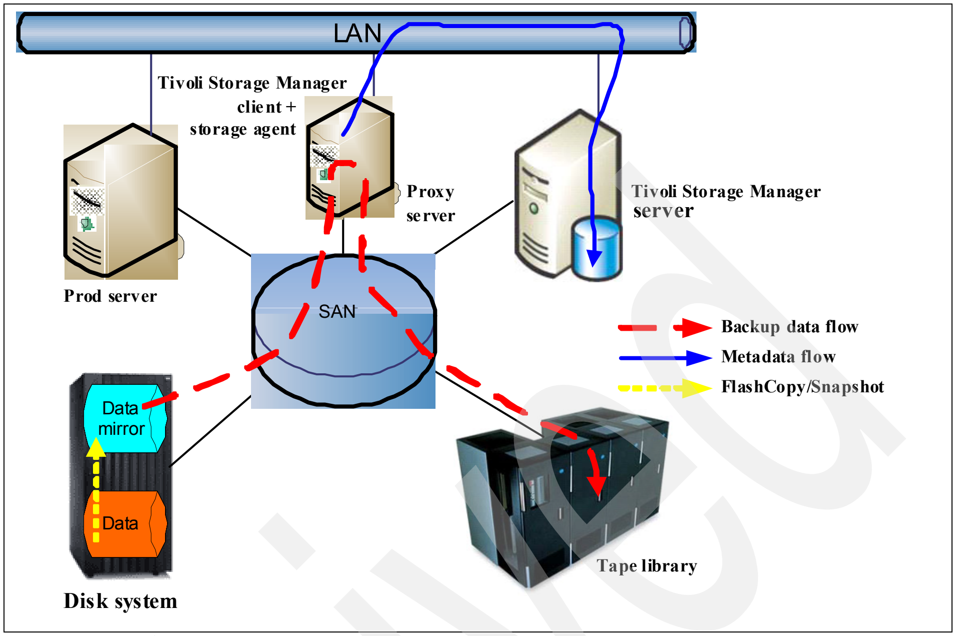

Figure 2-4   Split-mirror backup

2.2.5 Using disk systems for storing data

Tivoli Storage Manager has long supported using disk as a backup medium. Disk was primarily used earlier for caching file-based backups before they were migrated to tape. In recent years more and more companies are using disk to store data for longer periods of time or even permanently.

Disk storage can be utilized in the following two ways:

- Caching, when data is initially sent to the Tivoli Storage Manager server. Data is later migrated to tape (within a few days or weeks).
- Long-term or permanent storage.

You can implement disk storage in the following two ways:

- Random access disk (device type = disk)
- Sequential-access disk (device type = file)

Random access disk storage pools with disk device type

Storage pools with a device type of disk (“disk storage pools” hereafter) are mainly used for storing data temporarily. Data is sent from Tivoli Storage Manager clients and stored initially in disk storage pools. Data resides there for a short period of time, after which data is migrated to other storage (such as tape).

We generally do not recommend storing data permanently in disk storage pools because data becomes fragmented with time. Fragmentation results in a lot of wasted space that cannot be reclaimed easily.

Best practices:

- Have enough disk storage pool space to store one night’s file system backup. Send large file backups directly to tape.
- Consider using the cache=yes parameter on disk storage pools to decrease restore time.
- If you migrate data to tape on a daily basis, consider using unmirrored disks for storage pools. Although this introduces some risk, it cuts disk costs by 50% while not raising the risk of data loss nearly as much.
- A disk storage pool must reside on a high-performance disk system, such as the IBM DS4000™ or IBM DS8000 with FC disk.
- Use a separate disk system as a disk buffer, so you do not have any performance interference with other workloads.
- Do not store backup data and production data on the same disk system.

Sequential disk storage pools with file device type

Storage pools with a device type of file (“file storage pools” hereafter) are a common choice for storing data long term or permanently. Because of lower cost per megabyte of disk, many Tivoli Storage Manager designs have added more disks to the solution. Many designers are interested in designing systems that use no tape at all and save all backed up data to disk.

File storage pools have many advantages compared to disk storage pools and are the strategic choice of Tivoli Storage Manager Development. Future enhancements to disk-based storage pools are expected to be made to the file device type.

2.2.6 Using Virtual Tape Libraries for storing data

Virtual Tape Libraries (VTLs), such as IBM TS7520, are the perfect combination of disk systems and tape libraries. Essentially the TS7520 is a disk system controlled by a built-in computer that emulates a tape library and tape drives. The emulation is transparent to Tivoli Storage Manager, which sees the TS7520 as a library with tapes and tape drives. The TS7520 has the fast access capabilities of normal disk systems and the easy configuration and use of a real tape library.

Compression is also a key feature of the TS7520 because it compresses incoming data streams just as a real tape drive does. Compression enables the TS7520 to store much more data than a regular disk system utilized for storing Tivoli Storage Manager data.

You must consider a few issues before implementing a TS7520. Generally implementing and maintaining a TS7520 is more costly than a comparable tape system. Also, the TS7520 cannot match the speed of several tape drives combined, so backing up large amounts of data can still be achieved much faster by backing up directly to tape.

We recommend using the TS7520 (and other VTLs) for specific needs in a Tivoli Storage Manager environment. The TS7520 excels in random access and thus in backing up and restoring many small files. which are the main purposes of VTLs.

For more information about the TS7520 and VTLs in general, refer to *IBM TS7520 Virtualization Engine: Planning, Implementation, and Usage Guide*, SG24-7520.

2.2.7 Using tape for storing data

Most Tivoli Storage Manager designs incorporate tape and tape drives as the ultimate long-term storage location for backups and archives. Storing data on tape is the traditional way to perform backups and definitely has some advantages compared to disk systems and VTLs. Tape is much faster than disk storage when storing large files and is still the most cost-effective choice.

The main disadvantages for tape is its inability to read and write multiple data streams from multiple clients at the same time. Furthermore, tape is a sequential media and thus has a very slow seek time for individual files.

2.2.8 Choosing the right storage

Choosing where to store your backups and archives is the decision that has the most impact on the cost of your Tivoli Storage Manager solution. The larger the solution is, the more cost impact. Hardware and power cooling costs tend to rise at a steady rate when more data is stored and consequently more hardware is introduced into the environment.

Figure 2-5 illustrates the relationship between costs for hardware, power and cooling, and administration and the number of TBs stored in a Tivoli Storage Manager environment that retains its degree of complexity as the amount of data increases.

Choosing the right storage also involves choosing the right mixture of storage. The more types of storage you use, the more complex your environment becomes. This affects administration costs, so you need to find a balance that takes into account this factor.

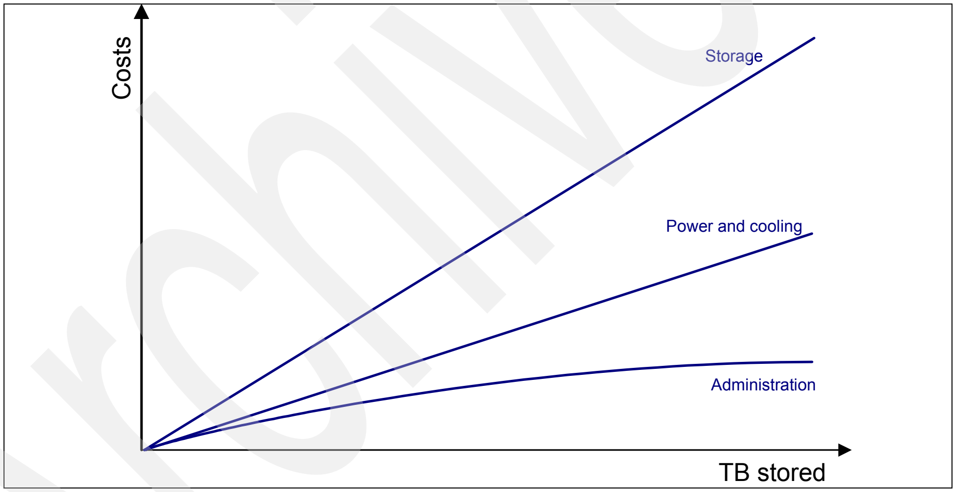

Figure 2-5 Relationship between cost and number of TBs stored

Comparing sequential disk, VTLs, and tape storage pools

When deciding whether to use sequential disk, VTLs, or tape as your primary long-term backup and archive storage, consider the following issues:

- Perform a realistic analysis of the total amount of storage you require using your data retention policies, an evaluation of the total amount of data on your clients, and your growth expectations. Evaluate the cost, functionality, and performance of the various storage types.
- Make sure you get the performance you need for the right backup and archive data. You do not have to buy high-performance storage for data that does not require it. Use your SLAs and the input you have gathered regarding backup and restore requirements.
- Make sure that you invest in a disk technology that can permanently store the data. Due to the larger risk of disk failure compared to tape, you must consider a copy storage pool for data redundancy (disk or tape pool). Serial Advanced Technology Attachment (SATA) disk systems are most commonly used for storing long-term data. These systems are, however, more sensitive to disk failures, which impact performance when rebuilding Redundant Array of Independent Disks (RAID), and generally provide poor performance compared to SCSI or Fibre Channel disk systems.
- Carefully consider the cost and technical feasibility of getting the data off site for disaster recovery purposes. Tape cartridges are portable. If you use sequential disk or VTLs, you must move the data using a network transport every night. How much does that transport cost and can it accommodate the nightly data workload? Consider using tape for generating off-site disaster recovery media.
- Sequential disk and VTLs provide fast access times to back up objects and thus faster restores, but carefully investigate where the restore performance bottlenecks lie before deciding to use sequential disk or VTLs as your primary storage pools. Bottlenecks can easily lie on the client system when restoring a large number of files (a common issue on Microsoft Windows systems). Perform extensive testing.
- Carry out correct calculations when examining how much data can be stored on a disk system. Very few systems provide compression (only some VTL systems do). This factor is a huge consideration when calculating possible utilization. Other disk-specific factors are the loss of usable disks due to RAID and OS formatting.

See Table 2-4 for a comparison of sequential disk, VTL, and tape storage pools.

Table 2-4 Benefits and drawbacks of sequential disk, VTL, and tape pools

Function	Sequential disk pool	VTL pool	Tape pool
Cost	► Arguably low acquisition costs on SATA disk systems ► High power consumption and cooling requirements ► Higher management costs	► Arguably low acquisition costs on most VTL systems ► High power consumption and cooling requirements ► Low management costs	► High acquisition costs for enterprise-type tape solutions ► Low ongoing costs as the environment grows ► Very low power and cooling costs
Backup performance	► Only single stream access to each virtual volume during backup ► Medium throughput for single backups ► I/O congestion on low-end SATA disk systems when doing multiple backups	► Only single stream access to each virtual volume during backup ► Medium throughput for single backups ► Intelligent disk layout and cache ensures constant backup performance	► Only single stream access to each volume during backup ► Very fast streaming ► Backup of large objects is very fast ► Use multiple tape streams for increased backup performance
Restore performance	► Very fast access time to backup objects ► Multiple stream access to each virtual volume during restore ► High throughput for single client restores with many small files ► I/O congestion on low-end SATA disk systems when doing multiple restores	► Very fast access time to backup objects ► Only single stream access to each virtual volume during restore ► High throughput for single client restores with many small files ► Intelligent disk layout and cache ensures constant restore performance	► Mount times heavily impact restore performance. ► Positioning for objects is relatively slow. ► Only single stream access to each volume during restore. ► Restore of large objects is very fast. ► Use multiple tape streams for increased restore performance (if used during backup).
Collocation requirements for fast restores	Preferred but not required	Not required but preferred, in some cases	Required for fast restore of many smaller files
Write validation	Some low-end SATA drives do not validate writes.	Write validation performed.	Write validation performed.

Function	Sequential disk pool	VTL pool	Tape pool
Redundancy requirements	► High failure rate of individual disks. ► RAID, mirroring, and copy pool generation is highly recommended.	► High failure rate of individual disks. ► RAID, mirroring, and copy pool generation is highly recommended.	Very low failure rate.
LAN-free capability	Complicated; possible only using IBM SANergy®	► Possible ► Medium performance	► Possible ► High performance
Off-site storage	Not feasible	Not feasible	Removable media
Hardware compression	► Not possible. ► Client software compression can be used as an alternative.	► Possible but not recommended because it decreases throughput. ► Client software compression can be used as an alternative.	Native.

Data placement recommendations

In concluding our comparison of the benefits and drawbacks of sequential disk, VTLs, and tape storage, this section discusses best practices for data placement.

Storing large backup objects

Figure 2-6 on page 62 shows where large objects, such as database objects and image backups of whole file systems, must be stored.

When you are backing up large objects and fast backup and restore is a requirement, we recommend you back up data directly to tape. Tape is, as mentioned previously, the fastest streaming device in your Tivoli Storage Manager environment. If your client can deliver the data fast enough, you can even use multiplexing and write to multiple tape drives at the same time. You can restore the data quickly and do so from multiple tape drives as well.

When you back up large files over slower connections, you must initially store the files in a disk storage pool because a slow running backup cannot deliver enough data to saturate a modern tape drive. You must therefore send multiple, slower backups to disk first and migrate them later to tape storage. Doing so ensures maximum utilization of your tape drives and thus reduces the Total Cost of Ownership (TCO) of your Tivoli Storage Manager environment. You can quickly restore data from either disk or tape.

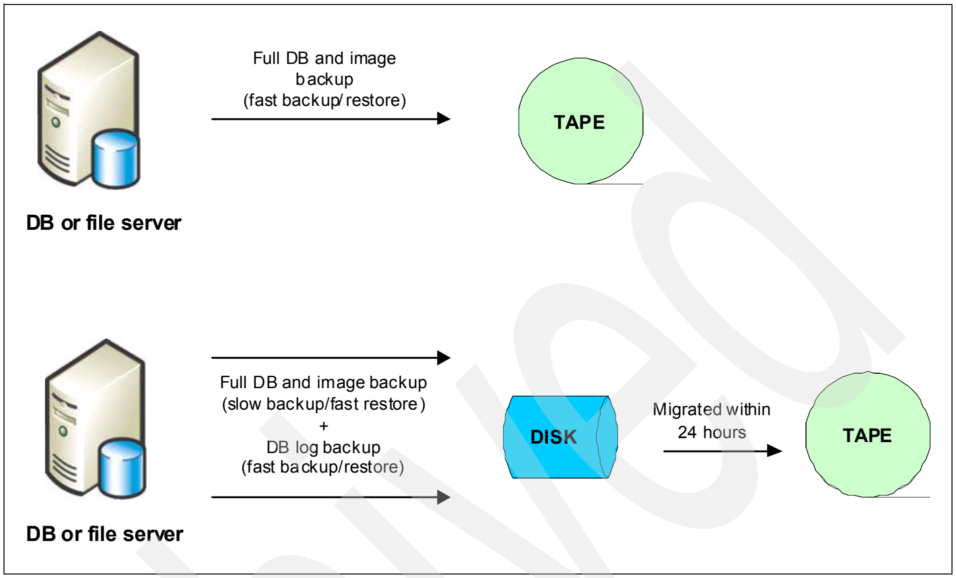

Figure 2-6 Data placement for large backup objects

Storing smaller backup objects

Smaller backup objects neither back up nor restore well from tape storage. The seek time of tape drives is too poor to get adequate performance and utilize them properly.

Figure 2-7 on page 63 depicts our recommendations for storing smaller backup objects. Critical data, which must be backed up and restored very fast, must be stored on storage with a fast access and seek time for individual objects. We recommend storing critical data permanently on file storage pools (sequential disk) or in VTLs.

Not all data is critical when it comes to backup and restore performance. Such data can easily be kept permanently on tape storage, but you must initially back it up to a disk-based device. Keeping the data on a disk-based device for a short period of time (for example, a week) dictates that you use disk storage pools, which are easiest to handle on a daily basis. Keeping the data on disk for 1-4 weeks requires a file of VTL-based storage pool because fragmentation plays a large role when the storage pool is disk based.

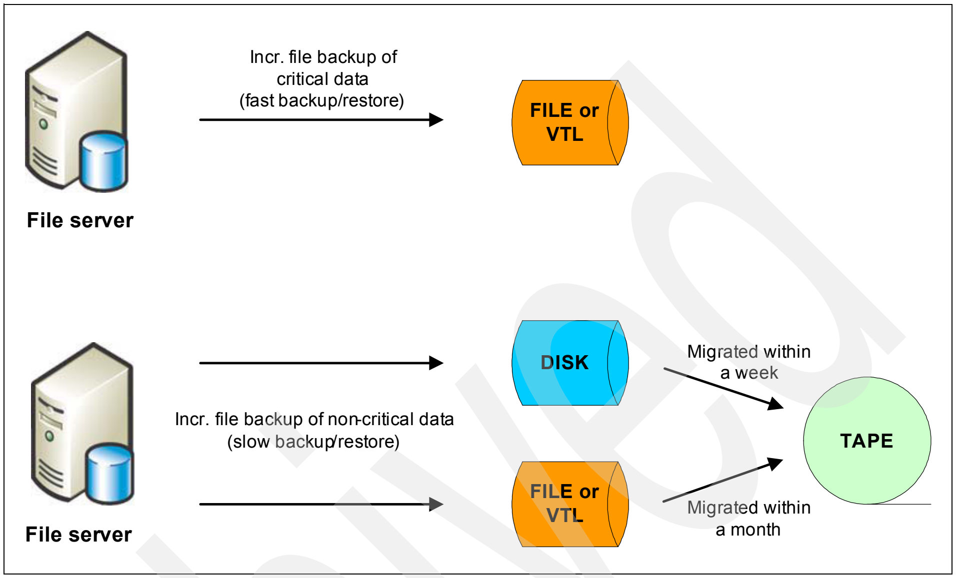

Figure 2-7 Data placement for small backup objects

2.3 Sizing your Tivoli Storage Manager solution

When sizing your Tivoli Storage Manager solution, you must determine the amount of storage you need (which are determined by retention policies), as well considering the hardware required to satisfy your SLAs (backup and restore performance requirements). The sizing process also involves a lot of calculations. In this section, we analyze the database and storage sizing.

2.3.1 Sizing considerations

This section describes the several factors you must consider when sizing a Tivoli Storage Manager solution.

Amount of data transported by the solution

To effectively design a Tivoli Storage Manager solution, you must understand how much data is transported by the solution. If the solution is doing backups, you must be able to move the data from the client nodes to the server on a nightly basis; therefore, all the equipment and infrastructure must be sized accordingly.

Most importantly, you must be able to restore any, or perhaps all, servers in the environment within an acceptable time frame.

The total amount of data that must be backed up and restored is a critical factor in the design because the Tivoli Storage Manager server, LAN or WAN transports, disk, and tape systems must all be sized to handle the amount of data that is backed up and restored. Typically, the cost of these items is the most expensive part of the Tivoli Storage Manager solution. It is critical that they be sized in a manner that minimizes expense while meeting backup and recovery requirements.

Types of data supported by the solution

The type of data you want to back up is related to the amount of data that is backed up and restored. Certain types of data, such as file system data, make use of an incremental forever strategy, which helps to minimize the amount of data that must be moved during backup. Other types of data, for example database and mail data, typically require periodic full backups, and many organizations perform a full backup each time to allow for quicker restores. Application servers sometimes contain millions of small files (<1 KB), which have a direct impact on the size of the Tivoli Storage Manager database and therefore on the number of required Tivoli Storage Manager server instances.

Retention policies

The Tivoli Storage Manager server must store data that is backed up and archived. The length of time an organization chooses to store this data directly affects the total amount of disk- and tape-based storage that is necessary in order to retain it and, of course, the cost of the solution as well. The retention policies you set in Tivoli Storage Manager determine the necessary amount of storage. Typically, retention policies are used to determine the amount of disk space or the size of the tape library that is required.

Data growth

When designing a Tivoli Storage Manager solution, you must consider data growth. Although a newly designed Tivoli Storage Manager solution might be appropriate for supporting your current IT environment, forgetting to plan for data growth can quickly cripple the solution and cause it to become unmanageable. In some cases, the amount of data in an environment can double in a year or less, and unless you account for this data in the initial solution design, the increase can cause incomplete backups and slow restores, ultimately resulting in a loss of data and in not meeting SLA requirements.

Small file performance

Systems with many small files (for example, 100 GB of 100 KB files) can be a backup system's worst problem. All storage devices perform at their worst when reading and writing small blocks of data compared to reading and writing large blocks of data. As you are developing SLAs for systems, you must account for such issues to avoid overcommitting on backup or restore times.

Best practices:

- When agreeing on SLAs, remember that restore performance is often dependent on the tape technology (which affects the seek time) and the speed the client can write files to the disk.
- For systems with many small files, we strongly recommend you consider the following alternate storage and backup methodologies to improve performance:
  - Sequential disk storage pools
  - Image backup
  - Collocation by group, node, or file space

Remember that when using its base functionality of incremental backup and file-level restore, Tivoli Storage Manager is simply making calls to the OS to enumerate the files for backup and then write them for restore. Although the Tivoli Storage Manager client might be somewhat faster due to multithreading, it cannot be "orders of magnitude" faster than the OS. Test-copy a set of files to and from the client disk using native OS utilities such as cp, tar, or Microsoft Backup. Use the information you learn in the test to establish a baseline for Tivoli Storage Manager.

Many times the bottleneck with small file operations is reading from or writing to the client disk. The only way to improve performance here is to improve client disk performance and not to change hardware on the backup system, for example, changing from storing data on tape to storing data on disk.

Tivoli Storage Manager has features that significantly improve small file restore performance. Specifically, assuming that client disk write operations are an issue, the IBM Tivoli Storage Manager image backup feature can usually restore a full disk drive much faster than file-by-file restore. The downside to using this feature is that you must restore the whole disk and cannot restore a single file to the client. When designing the solution for a system with many small files, architects usually choose to combine image backup with incremental backup to achieve quick restore in case of total loss and file-level granularity.

Collocation

Finally, we must note that the restore speed of many small files is only as good as the organization of data on the tapes (assuming that you are using tape backup). The default for Tivoli Storage Manager storage pools is to use group collocation. If this default is changed to no collocation, Tivoli Storage Manager comingles data from multiple systems on the same volume. When using an incremental forever methodology, it can cause backup data to become fragmented across all volumes in the storage pool, which in turn can increase restore times. Collocation causes backup data from a specific system or set of systems or file space to be kept together, on as few volumes as possible, thus decreasing fragmentation and improving restore times.

Beginning with Tivoli Storage Manager V5.3, three types of collocation are available: collocation by group, collocation by node, and collocation by file space. Each of these methodologies for grouping files, along with its pros and cons, is described in Table 2-5.

Table 2-5 Collocation types

Collocation type	Description	Pros	Cons
None	Files backed up from multiple nodes are comingled on all tapes. The location is completely dependent on when the data was written. Tivoli Storage Manager optimizes the writing of data to tape to maximize the amount of data stored on a tape and the time it takes to write the data.	► Best possible times when moving data to tape ► Most efficient use of tape space ► Fewest tape mounts during backup or housekeeping jobs	► Trades restore efficiency for backup efficiency. ► Might require numerous mounts and seeks to restore an entire system of small files. ► Restoring clients using the same tapes is not possible to do at the same time.
Collocation by group	Tivoli Storage Manager default. Nodes are grouped on tape in user-defined groups. Groups can be defined to reduce tape waste while still preventing data from becoming too fragmented.	► Reduces restore times by decreasing data fragmentation and mounts ► Still allows for efficient use of tape ► Reasonable number of tape mounts during a backup or housekeeping job	► Requires figuring out optimum node-to-group assignments. ► Can cause some tape waste. ► Increases backup or migration time (or both) due to increased mounts. ► Restoring clients using the same tapes is not possible to do at the same time.

Collocation type	Description	Pros	Cons
Collocation by node	Each node is assigned its own tape or a set of tapes.	Reduces restore times by decreasing data fragmentation and mounts	► Can waste a significant amount of tape (and library slots) if nodes are considerably smaller than tape size. ► Increases backup or migration time (or both) due to increased mounts. ► Can work against multisession restore.
Collocation by file space	Each file space is assigned its own tape or a set of tapes.	Reduces restore times by decreasing data fragmentation and mounts	► Can waste a significant amount of tape (and library slots) if file spaces are considerably smaller than tape size. ► Increases backup or migration time (or both) due to increased mounts. ► Can work against multisession restore.

Which version of collocation is the right one to use? No universal response can answer this question. The restore times associated with each version have a different cost in terms of hardware and benefit. As a general rule, when storing small files on tape, you do not want them spread out over too many tapes. Having them spread out like that usually leads to some sort of collocation by group for most systems where you allocate enough systems to the group to fill up one or more tapes. You might also use collocation by file space or node for systems that have large numbers of small files and large amounts of data (large file servers). Database, mail, and ERP clients usually benefit equally from both collocation and no collocation. They store few but very large objects on the Tivoli Storage Manager server and therefore fragmentation has very little impact on restore performance.

If in doubt, use collocation by group on your storage pools, both primary and copy. When you devote time to figuring out the optimum node-to-group assignment, the space wasted on tapes can be negligible compared to the benefits of the simplicity of the design. Group small nodes together to keep them on fewer tapes and leave your bigger clients ungrouped, because they then are automatically collocated by node. So even if you forget to put a node in a collocation group and the node makes use of a storage pool with group collocation, the worst that can occur is the node uses a whole tape by itself. It never becomes uncollocated using this strategy.

Sequential disk storage pools and tape storage pools make use of the same collocation rules.

Note also that the same collocation methodology does not necessarily have to be used for all nodes on a Tivoli Storage Manager server. Some more important ones might ultimately store their data in collocated pools, and less important ones might forgo collocation to reduce backup times and save money on tapes and tape drives.

2.3.2 Backup workload calculations

An important consideration when designing the overall Tivoli Storage Manager solution is the total amount of data that the Tivoli Storage Manager server requires to support backup or restore over a particular time frame.

Best practices:

- Calculate the amount of data the Tivoli Storage Manager server has to accommodate during its nightly backup window and ensure that the hardware can reasonably accommodate the data.
- Calculate the amount of data that is restored during disaster recovery and ensure that the hardware meets disaster recovery SLAs.

For example, we can calculate the workload for a Tivoli Storage Manager server in a theoretical environment. The environment has the following systems that we must back up:

- 20 file servers
  - Each server has 60 GB of file data.
  - Incremental forever backup with a change of approximately 10%.
  - Six hour backup window.
- Four DB servers
  - Each server has 100 GB of database data.
  - Full backups with a change of 100%.
  - Six hour backup window.

So doing the calculation for backup workload, each night we must back up:

$$(60 \times 20 \times .10) + (100 \times 4 \times 1) = 520GB$$

And we divide that by the time allowed for the backup:

$$\frac{520GB}{6Hr} \approx 86.6 \frac{GB}{Hr}$$

So the network interfaces, tape drives, Tivoli Storage Manager server, and all other infrastructure must be able to handle this amount of data in order to meet SLAs for this environment.

Note: This type of calculation does not apply only to backups. It can and must be performed for disaster recovery restores and single system restores.

2.3.3 IBM Tivoli Storage Manager database

The Tivoli Storage Manager database is critical to the operation of the entire Tivoli Storage Manager environment, and you must exercise great care when designing the solution. When planning a Tivoli Storage Manager environment, you must make sure the database is kept below a reasonable size and implement more than one Tivoli Storage Manager instance if required, to keep the database sizes within limit.

Best practices:

- Do not permit the Tivoli Storage Manager database to get too large. A good general rule is 120 GB, but no “magic” number applies. When expiration, database restores, and other Tivoli Storage Manager administration processes take too long and client restores become too slow, the database is too large.
- When the database reaches 80 GB, we recommend you create no additional clients on that Tivoli Storage Manager instance. Databases tend to grow over time due to fragmentation.
- Consider using Tivoli Storage Manager internal software mirroring. It is the fastest and most reliable way to perform software mirroring.
- Spread Tivoli Storage Manager volumes over as many physical disks as possible for better performance.
- Use smaller DB volume sizes (for example, 4 to 8 GB) because they improve database performance. Keep the number of database volumes below 20.

See 6.2.2, “Infrastructure performance tuning” on page 279 for further discussion on database performance.

You can use several methods to predict the size of the Tivoli Storage Manager database. The first and the easiest method is this general rule: Given an environment consisting of mixed size files, a typical Tivoli Storage Manager database is 3-5% of the total amount of data in the environment. For example, suppose you are backing up 20 servers with 60 GB of file data each, or 1200 GB of data; the Tivoli Storage Manager database is approximately 36-60 GB.

Another way to estimate the size of the Tivoli Storage Manager database is by using the total amount of data and average file size. Table 2-6 summarizes the amount of space each file or object uses in the database.

Table 2-6 Space occupied in database

Description	Size
Initial file/object backup	600 bytes
Each additional version	200 bytes
Each copy storage pool copy	200 bytes

Using the information from the previous example, assume you back up 20 servers with 60 GB of file data each, and that you have one copy pool and wish to keep 15 versions of the file. Also assume that the files on these systems are small, having an average file size of 100 KB each. We can calculate the estimated size of the Tivoli Storage Manager database as follows:

- Total data to back up is 20 servers x 60 GB = 1200 GB.
- Number of files 1200 GB / 100 KB = 12.3 million files.
- With 1 copy pool and 15 versions, each files uses:
  600 bytes + (14 copies x 200 bytes) + (15 storage pool copies x 200 bytes) or 6400 bytes.
- With 12.3 million files, you need 12.3 million x 6400 bytes or 73.3 GB.

2.3.4 Tape drives

It is important to understand the performance characteristics of tape drives and tapes when designing a Tivoli Storage Manager solution.

Best practices:

- When using tape storage pools, a Tivoli Storage Manager server must have access to no less than two drives.
- For architecture calculation purposes, assume only 80% of maximum uncompressed throughput for a tape drive.
- Be prepared to perform restores while other administrative operations are happening on the system or when drives are broken.
- Carefully consider card and bus throughput when attaching tape drives to systems. Most protocol and tape combinations can accommodate two or three tape drives per card.
- When using Fibre Channel/SAN attached tape drives and disks, do not mix disk and tape traffic on the same host bus adapter (HBA).

When planning for tape drive throughput and tape capacity, it is best to be conservative. Although drives can perform close to and better than their compressed rating, their performance is typically far less than their rating in reality. We use 80% of native uncompressed throughput ratings and 150% of native capacity ratings for theoretical calculations of drive throughput and tape capacity. Table 2-7 summarizes those calculations for some popular tape drives.

Table 2-7 Tape drive calculations

Tape drive	Native speed (MBps)	Native capacity (GB)	Assumed speed (GB/hr)	Assumed capacity (GB)
IBM LTO Gen 1	15	100	42	150
IBM LTO Gen 2	35	200	98	300
IBM LTO Gen 3	80	400	224	600
IBM LTO Gen 4	120	800	338	1,200
IBM 3590-E1A	14	40	39	60
IBM 3590-H1A	14	60	39	90
IBM 3592 J1A	40	300	112	450
IBM 3592 TS1120	104	500	298	750

Tape drive	Native speed (MBps)	Native capacity (GB)	Assumed speed (GB/hr)	Assumed capacity (GB)
SUN T9840C	30	40	84	60
SUN T9940C	30	200	84	300
SUN T10000	120	500	298	750
Sony SAIT	30	500	84	750

Whether you use Table 2-7 on page 71 or your own calculations, it is critical when designing a Tivoli Storage Manager solution to measure the backup and potential restore workload against the speed of your tape solution. In our earlier example, for instance, we calculated that we needed to move 86.6 GB per hour to meet SLAs. So based on the theoretical numbers in Table 2-7 on page 71, a solution that is capable of that data movement requires three LTO Gen 1 tape drives.

In addition to calculating backup workload, an architect must also consider administrative processes, potential restores, and broken drives when deciding the number of tape drives a system requires. Each day, off-site copies of tapes must be made, and reclamation must occur. A potential restore of a large system can also require multiple drives to meet SLAs.

A detailed discussion of specific I/O protocol speeds is beyond the scope of this book. However, you must understand the throughput speeds and seek times of the drive you choose and recommendations for system connectivity. As a general rule, most hardware vendors recommend no more than two to three tape drives per interface card, and if using Fibre Channel, this tape traffic must be separated from disk traffic.

Important: When doing restores of millions of small files, the seek time of the drive can directly dictate the time required to perform the restore, and native drive throughput becomes a negligible factor.

2.3.5 Tape libraries

Tivoli Storage Manager was designed with automation in mind, specifically the use of automated tape libraries. The best and easiest way to run a Tivoli Storage Manager library is to keep all on-site data in the library so that tapes can be mounted automatically when required for restores, backups, reclamation, and other Tivoli Storage Manager administration processes.

Best practices:

- Size your tape library so that it can accommodate all of your on-site data: Count on seven database backups, scratch tapes, and 6 to 12 months of growth.
- When calculating library capacity, assume 1.5x data compression and 65% tape utilization.
- Choose a library with sufficient mount speed. If you are using collocation on tape pools, a mount is required for each collocated client or collocation group multiplied by the number of tape storage pools that the data is to be written to.

For example, suppose we are calculating the number of slots required in a library that uses the data from the workload calculations with some additional information about retention policies and growth.

The environment has the following systems that we must back up:

- 20 file servers
  - Each server has 60 GB of file data.
  - Incremental forever backup with a change of approximately 10%.
  - Six hour backup window.
- Four DB servers
  - Each server has 100 GB of database data.
  - Full backups with a change of 100%.
  - Six hour backup window.

Data must be kept for 32 days. Growth in the environment is approximately 10% per year. To compute total data storage (assuming no on-site copy storage pool), we use the following multiplication:

{(20x60x1)+(20x60x.1x31)}+(32x4x100)

{1200+3720}+12800=17720

For this example that we use LTO Gen 2 tapes. Using the 65% rule and 1.5x compression, we can assume that tapes in the library hold an average of 195 GB. This results in:

17720/195 = 91 tapes

Approximately 91 tapes are required. Adding room for 10% growth, 7 database backups, and 10 scratch tapes, we can choose a library with 117 or more slots.

It is best that you choose a library that you can easily expand so that future data growth can be accommodated without a complete library replacement.

2.3.6 Sequential disk, VTL, and tape utilization comparison

Calculations on utilization are important when deciding what technology to use for long-term storage of your backup and archive data. When planning your Tivoli Storage Manager environment, you need to know actual client data you can store on a specific storage type because this information helps you decide what kind of storage you can acquire that is within your budget. As a basic rule, you cannot assume that all the GBs of disk or tape you plan to purchase can be used to store your backup data. Different technologies have different utilization percentages.

In this section, we consider disk-, VTL-, and tape-based examples. The purpose of the examples is to discover which parameters you must take into consideration when doing your calculations. The examples are meant as guidelines; you have to carry out your own calculations taking into account the kind of hardware you have, your retention requirements, and solution scope.

The calculations only address raw capacity and utilization. No costs are being factored in at this point.

Before doing the actual calculations for the three types of storage, we take a look at the reasoning behind the calculations. Each calculation is based on a list of assumptions. These assumptions detail the initial amount of storage used for the calculations as well as all the factors influencing and subtracting storage space (space that cannot be used for storing TSM client data). When subtracting this storage space from the initial amount of storage, we eventually end up with the amount of storage that can actually be used by Tivoli Storage Manager to store client data, and we are also able to calculate a utilization percentage of the initial storage.

Each calculation table in 2.3.1, "Sizing considerations" on page 63 lists the initial amount of raw storage capacity at the top (112 TB in each case). From this calculation several factors are subtracted:

- Spare volumes used in disk systems
- System configuration taking up space in disk systems
- RAID formatting taking up space in disk systems
- Filling Tivoli Storage Manager volumes with almost no data on them
- Storage used for Tivoli Storage Manager reuse delay
- Storage wasted in Tivoli Storage Manager file aggregates

During the calculation, subtotals are listed if relevant for the specific calculation:

- Usable before raid formatting
- Usable by Tivoli Storage Manager
- Usable for client data

At the end of each table, the result of the entire calculation is listed as “Total GB client data that can be stored.”

Sequential disk system

We make the following calculations for disk utilization:

- An IBM DS4800 equipped with SATA disks is used in this example.
- The DS4800 disk system consists of 14 EXP810 enclosures, each containing 16 disks totaling 224 physical disks in all. A physical disk is 500 GB in size giving a total capacity of 112 TB.
- For every second disk enclosure, a spare volume is set aside, eight spare volumes in all.
- From each physical disk, 7% is subtracted. This space is used by the storage system to retain information about the controller and the disks’ placement within the disk system.
- The physical disks are divided into array-groups, which are RAID 5 formatted. This results in approximately 14% of the remaining GBs being used for parity information.
- The individual Tivoli Storage Manager file volumes on the disk system are predefined and set to a size of 50 GB.
- The requirement for scratch volumes is not addressed. The calculations are mainly focused on raw capacity and theoretical utilization.
- Group collocation is used in the storage hierarchy, of which there is only one. The Tivoli Storage Manager solution consists of 100 clients divided evenly into 10 collocation groups. Another 20 clients are not in collocation groups and therefore are collocated by node. Each collocation group or collocated node, in its worst case, has one volume each in filling mode with 1 KB of data on it. Thirty volumes have to be allocated for this use in order to have enough capacity. For a mere calculation on basic capacity and utilization, these 30 volumes can be disregarded. However, the collocation factor has been included to show the difference between sequential disk and tape, when using the same type of collocation, because we recommend using 20-50 GB volumes and some type of collocation. This configuration provides the best performance without impacting utilization of the disk system too heavily.
- The utilization of file volumes on the disks are not, on average, 100% due to the requirement to keep the reclamation percentage at approximately 20%.

This requirement is necessary because I/O on the disk system, from reclamation processes, has to be kept at a certain level. Also more reclamation results in more pending volumes due to reuse delay, which in turn results in more unusable space for “active” data. In general, 10% space is wasted by inactive file aggregates when setting the reclaim percentage to 20%.

- We estimate 1% of daily changes in the Tivoli Storage Manager environment each day, meaning 1% of the total amount of stored data is changed every day. We also want to be able to restore the Tivoli Storage Manager server 3 days back without loss of data and without having to audit all volumes (reuse delay = 3). In general, this means it takes 20 days before a volume is reclaimed (due to the reclaim percentage of 20%). The volume is then reclaimed, and 3 days after the reclamation occurs, the volume is put back into circulation. Because we have a 23-day cycle, the volume is unavailable 3/23 days or 13% of the time due to reuse delay. This estimate is, however, a rough estimate, but it is useful for making the point that low reclamation values produce many volumes eligible for being retained from use by the reuse delay factor.

Table 2-8 shows that only approximately 65 TB out of 112 TB raw disk capacity can be utilized for Tivoli Storage Manager data (before the you need to add more capacity). When using these numbers to calculate the utilization precentage of the initial storage, we arrive at 58% (65/112 * 100).

Table 2-8 DS4800 calculation

Description	**Available storage**
DS4800 raw capacity (224 disks x 500 GB)	112,000 GB
- 8 spare volumes	- 4,000 GB
- 7% for system configuration	- 7,560 GB
Usable before raid formatting	100,440 GB
- 14% for RAID5 formatting	- 14,062 GB
Usable by TSM	86,378 GB
- 30 filling volumes (30 volumes x 50 GB)	- 1,500 GB
Usable for client data	84,878 GB
- 13% for reuse delay	- 11,034 GB
- 10% for wasted space in file aggregates	- 8,488 GB
Total GB client data that can be stored	**65,356 GB**

VTL system

We make the following calculations for VTL system utilization:

- An IBM TS7520 Virtualization Engine™ is used for the calculations.
- The TS7520 has fourteen SC6/SX6 cache modules. Each SC6/SX6 contains sixteen 500 GB disk drives, giving a total of 8 TB of unformatted storage per SC6/SX6 and grand total of 112 TB for the whole TS7520.
- The disk arrays reside in a RAID 5 stripe, with parity and hot spares using both seven data plus one parity (7+P) and six data plus one parity and one spare (6+P+S). For RAID 5 arrays, one disk worth of parity is spread across all disks in the array. One disk in each TS7520 Cache module is designated as a spare. This gives a usable capacity of 6.5 TB per SC6/SX6.
- Fifty GB virtual tapes are defined for the TS7520. No volumes are preallocated.
- The requirement for scratch volumes is not addressed. The calculations are mainly focused on raw capacity and theoretical utilization.
- Group collocation is used in the storage hierarchy, of which there is only one. The Tivoli Storage Manager solution consists of 100 clients divided evenly into 10 collocation groups. Another 20 clients are not in collocation groups and therefore are collocated by node. Each collocation group or collocated node, in its worst case, has one volume each in filling mode with 1 KB of data on it. Because no volumes are preallocated, the storage the filling volumes use can be disregarded. Collocation is used, so clients' data is spread across as few virtual tapes as possible to facilitate noncongested restores, when restoring many clients at the same time.
- The utilization of the virtual tapes are not, on average, 100% due to the requirement to keep the reclamation percentage at approximately 20%. This percentage is required because I/O on the disk system, from reclamation processes, has to be kept at a certain level. Also more reclamation results in more pending volumes due to reuse delay, which in turn results in more unusable space for "active" data. In general, 10% space is wasted by inactive file aggregates, when setting the reclaim percentage to 20%.
- We estimate 1% of daily changes in the Tivoli Storage Manager environment each day, meaning 1% of the total amount of stored data is changed every day. We also want to be able to restore the Tivoli Storage Manager server three days back without loss of data and without having to audit all volumes (reuse delay = 3). In general, this means it takes 20 days before a volume is reclaimed (due to the reclaim percentage of 20%). The volume is then reclaimed, and three days after the reclamation occurs, the volume is put back into circulation. Because it is a 23-day cycle, the volume is unavailable 3/23 days or 13% of the time due to reuse delay. This is, however, a rough

estimate but useful for making the point that low reclamation values produce many volumes eligible for being retained from use by the reuse delay factor.

Table 2-9 shows that only approximately 71 TB out of 112 TB raw tape capacity can be utilized for Tivoli Storage Manager data (before the requirement arises to add more capacity). When using these numbers to calculate the utilization precentage of the initial storage, we arrive at 64% (71/112 * 100).

Important: We have not taken compression into account when calculating TS7520 raw capacity and utilization. It is possible to gain a 150-200% increase in capacity with compression. Doing so, however, affects performance and might not be favorable.

Table 2-9 TS7520 calculation

Description	Available storage
TS7520 raw capacity (14 SC6/SX6 x 8 TB)	112,000 GB
- spare volumes and raid formatting (14 x 3 x 500 GB)	- 21,000 GB
Usable by TSM	91,000 GB
- 0 filling volumes	- 0 GB
Usable for client data	91,000 GB
- 13% for reuse delay	- 11,830 GB
- 10% for wasted space in file aggregates	- 9,100 GB
Total GB client data that can be stored	**71,253 GB**

Tape system

We make the following calculations of tape system utilization:

- An IBM TS1120 tape drive using 500 GB tapes is used for the calculations.
- A physical tape is 500 GB in size and 224 tapes are used to provide a capacity of 112 TB.
- The requirement for scratch volumes is not addressed in these calculations. The calculations focus mainly on raw capacity and theoretical utilization.
- Group collocation is used on the storage hierarchy, of which there is only one. The Tivoli Storage Manager solution consists of 100 clients divided evenly into 10 collocation groups. Another 20 clients are not in collocation groups and therefore are collocated by node. Each collocation group or collocated node, in its worst case, has one volume in filling mode, with 1 KB of data on it.

Thirty volumes have to be allocated for this use in order to have enough capacity. For a mere calculation on basic capacity and utilization, these 30 volumes can be disregarded. However, the collocation factor has been included to show the difference between sequential disk and tape backup, when using the same type of collocation because we recommend using at least group collocation on primary storage pools. Using collocation gives the best restore performance without impacting too heavy on utilization of the tape media.

- Reclamation percentage is set to 80%. This percentage is chosen due to the size of the tapes and the time required to reclaim individual tapes. Experience has shown that larger tapes take a considerable amount of time to reclaim. In general, 40% space is wasted by inactive file aggregates when setting the reclaim percentage to 80%.
- We estimate 1% of daily changes in the Tivoli Storage Manager environment each day, meaning 1% of the total amount of stored data is changed every day. We also want to be able to restore the Tivoli Storage Manager server three days back without loss of data and having to audit all volumes (reuse delay = 3). In general, it takes 80 days before a volume is reclaimed (due to the reclaim percentage of 80%). The volume is then reclaimed, and three days after the reclamation occurs, the volume is put back into circulation. Because we have an 83-day cycle, the volume is unavailable 3/83 days or 3.6% of the time due to reuse delay. This estimate is, however, a rough estimate but useful for making the point that higher reclamation values produce less volumes eligible for being retained from use by the reuse delay factor.

Table 2-10 shows that only approximately 55 TB out of 112 TB raw tape capacity can be utilized for Tivoli Storage Manager data (before the requirement arises to add more capacity). When using these numbers to calculate the utilization precentage of the initial storage, we arrive at 49% (55/112 * 100).

Important: No compression has been taken into account when calculating raw tape capacity and utilization. Generally, you must estimate the final capacity to be 150-250% of native capacity (depending greatly on the type of data stored on the media).

Table 2-10 TS1120 tape calculation

Description	Available storage
TS1120 tape raw capacity (224 tapes x 500 GB)	112,000 GB
Usable by TSM	112,000 GB
- 30 filling volumes (30 volumes x 500 GB)	- 15,000 GB

Description	Available storage
Usable for client data	97,000 GB
- 3.6% for reuse delay	- 3,492 GB
- 40% for wasted space in file aggregates	- 38,800 GB
Total GB client data that can be stored	**54,708 GB**

Conclusion about utilization comparison

The utilization of the three kinds of storage is somewhat similar ranging from 49% to 64%. The two disk-based solutions have almost the same utilization, with tape falling 10%-15% behind.

Taking hardware compression into account, the TS7520 and TS1120 can clearly store much more data than the native disk system. If considering maximum performance as a requirement as well, hardware compression is only applicable for the tape-based solution. This gives the TS1120 solution a 98% utilization compared to 58% and 65% for the DS4800 and TS7520, respectively!

Note: The value for TS7520 and hardware compression is in parenthesis in Table 2-11 on page 80 to indicate that compression is possible but not recommended.

Table 2-11 illustrates the differences.

Table 2-11 Utilization comparison

Hardware	Raw capacity (GB)	Utilization (%)	Utilization (GB)	Utilization with HW compression 2:1 (GB)
DS4800	112,000 GB	58%	65,356 GB	N/A
TS7520	112,000 GB	64%	71,253 GB	(142,506 GB)
TS1120	112,000 GB	49%	54,708 GB	109,416 GB

Figure 2-8 depicts the differences in utilization of the three storage types (with and without taking compression into account). The "Utilization w. compression" bar for TS7520 is shown in lighter colors to emphasize that although using compression with the TS7520 results in the best utilization of the three, it must be used with care and consideration because it decreases performance.

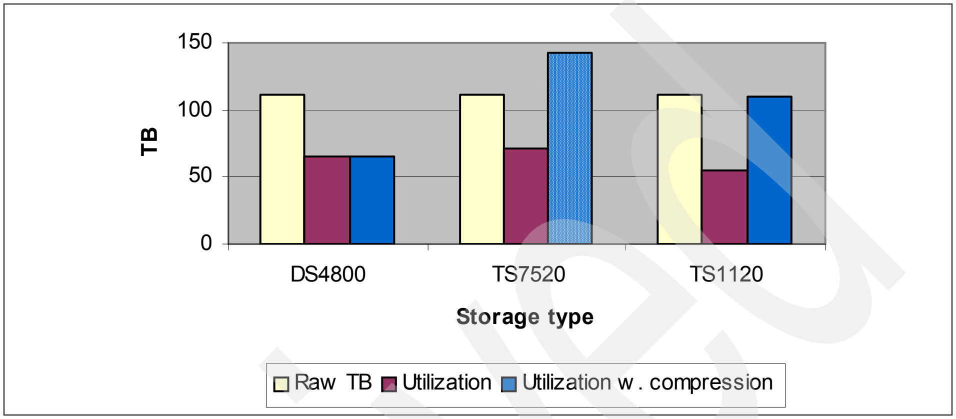

Figure 2-8 Illustrating utilization of DS4800, TS7520, and TS1120

We can conclude you must carefully plan and consider these issues before choosing hardware for your long-term data storage because your choice heavily impacts performance, as well as the costs of your Tivoli Storage Manager environment.

2.3.7 Sequential disk, VTL, and tape cost comparison

In the previous section, we calculated on utilization only. In this section, we factor in costs for hardware and power and cooling as well. These calculations address the true Total Cost of Ownership (TCO) and provide the base upon which we can decide what storage to use for storing our client data in the long term. These calculations are theoretically sound and indicate the costs of the hardware at the time of writing this book. Hardware prices change on a daily basis, and you must always do your own calculations.

To do ours, we first need to establish a baseline. We make the following assumptions:

- All power used to run hardware is converted into heat. This is true for 99% of the power used, but for calculation purposes we assume 100%.
- Cooling the hot air generated requires 100% more power to cool the air down again. Probably closer to 2-3 times the power is needed, but for calculation purposes we assume 100%.

- The baseline for the implementation is 112 TB with an annual growth of 25%.
- A kilowatt hour costs $0.30.
- Floor space and other data center infrastructure are not addressed in the calculations and of course must be included in a real life scenarios.
- An IBM TS3500 library with eight IBM TS1120 tape drives and 224 (500 GB) tape cartridges is configured and used as the tape baseline. Total cost for 112 TB raw capacity including monthly maintenance per year is $590,000. All tape hardware is replaced every sixth year.
- The TS3500 library draws 185 Watts (W), and the TS1120 tape drives draws 50 W each. This amounts to 585 W or 5125 kWh per year.
- Two IBM DS4800 disk systems, each with 8 GB cache and seven EXP810 expansion units with 500 GB SATA disks are used. Total cost is $880,000 for 112 TB of disk. Disk hardware is replaced every third year because maintenance is too costly after three years.
- The DS4800 controllers require 540 W each and the EXP810 expansion units 390 W each. This amounts to 6530 W or 57,203 kWh per year.
- One IBM TS7520 VTL system with one CV6 virtualization engine, four SV6 cache controllers, and ten SX6 Cache modules is used as the TS7520 baseline. Each SV6 and SX6 uses 500 GB SATA disks and has 8 TB raw capacity.
- The CV6 draws 510 W and the SV6/SX6 390 W each. This amounts to 5970 W or 52,297 kWh per year.

Power and cooling costs

Energy use is currently becoming more important, and managing energy efficiency is a major task in data centers. Let us take a look at the costs for powering and cooling 112 TB of different kinds of storage over a six-year period.

Figure 2-9 shows the amount in thousands ($) that is need to power and cool 112 TB of raw storage over a six-year period. As you can see, the costs are much bigger for disk-based systems than for the tape-based system, which of course is no big surprise given the previous assumptions. Each graph is growing exponentially due to the assumed 25% yearly growth in data.

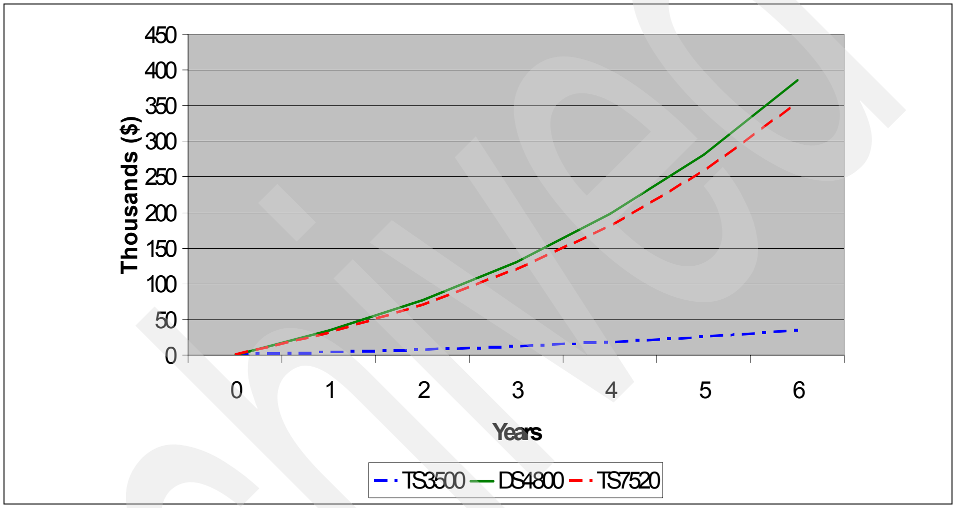

Figure 2-9 Power and cooling costs for 112 TB raw storage

Comparing raw storage capacity is however not completely correct. The comparison does not show the true costs of choosing and running specific hardware with Tivoli Storage Manager. We must take into account the utilization factor we calculated in 2.3.6, “Sequential disk, VTL, and tape utilization comparison” on page 74, as well as the ability to use hardware compression.

Figure 2-10 on page 84 shows the true cost of powering and cooling the three kinds of storage over a six-year period. The figure shows that disk-based solutions require even more power per stored TB client data than previous illustrated. This increase is a direct result of the inability to use hardware compression and of being able to store more backup client data. More than $600,000 can be saved on electricity during a six-year period if your data is stored on TS1120 instead of storing it on DS4800s.

Two graphs are shown for the IBM TS7520; one in which data is stored with hardware compression and one without.

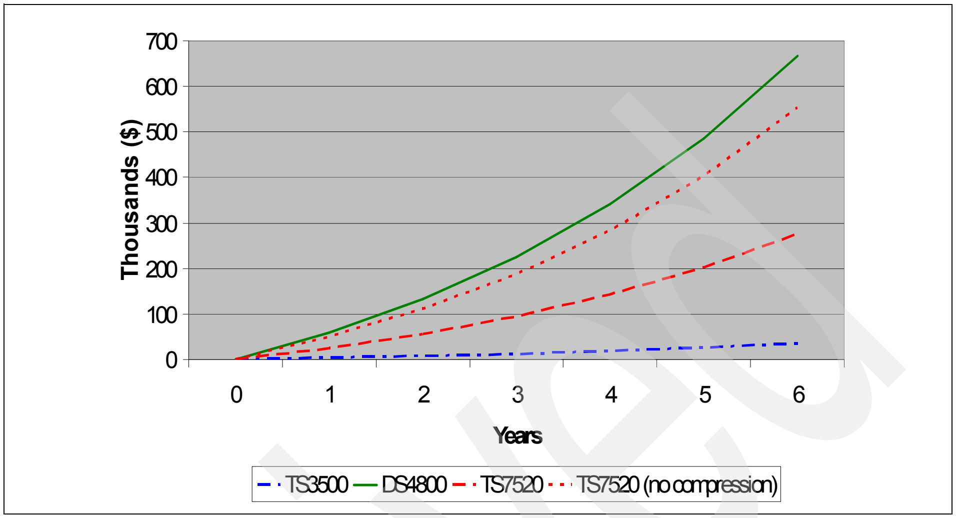

Figure 2-10 Power and cooling costs considering utilization and compression

Total Cost of Ownership

Besides power and cooling, we need to address the TCO of the three types of storage. TCO in this example includes hardware, maintenance, and power and cooling costs. We look at a six-year period.

Figure 2-11 shows the difference in TCO over a six-year period. The graph takes all issues previously addressed into consideration (hardware, maintenance, power and cooling costs, as well as possible utilization of storage and hardware compression). The graphs show that native disk systems are more than three times as expensive as tape-based solutions. Comparing VTL and tape gives tape a clear advantage again, being 2-4 times less costly.

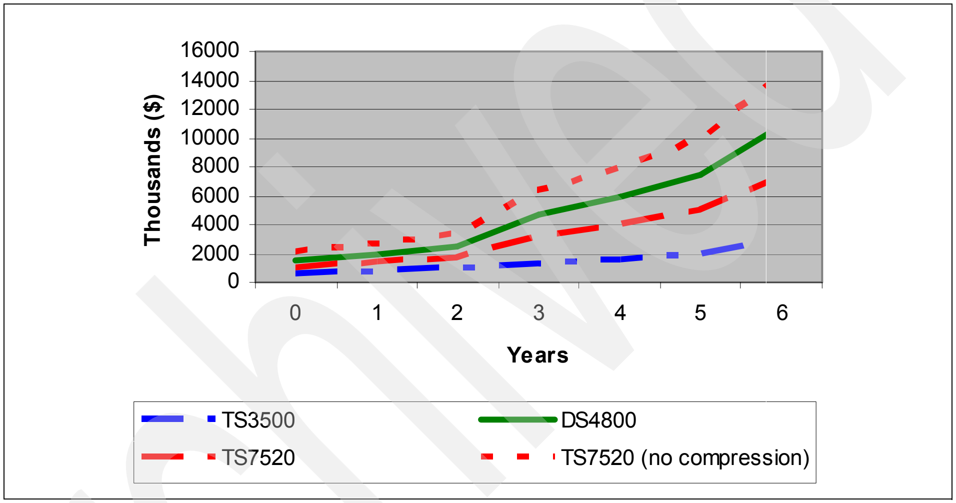

Figure 2-11 Six-year TCO analysis

Figure 2-12 illustrates the TCO after six years. As you can see, the main cost is for hardware and maintenance.

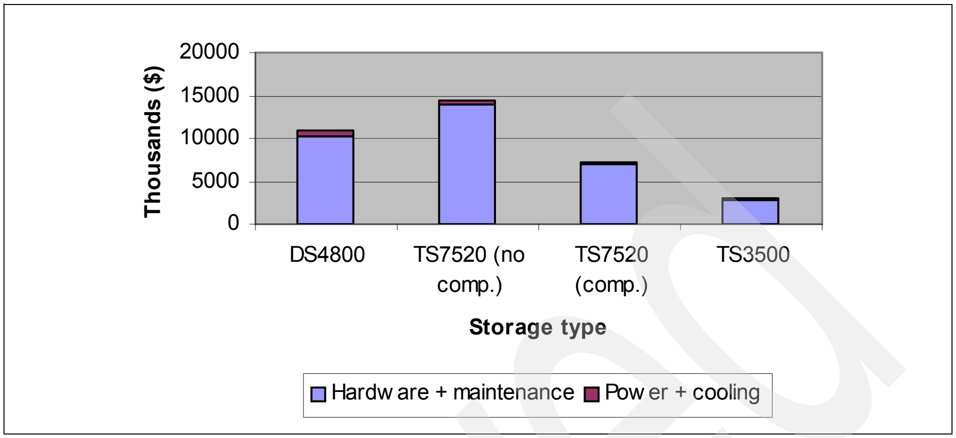

Figure 2-12 TCO after six years

Conclusion

Looking back at our assumptions and calculations, it is apparent that tape has a vast cost advantage compared to disk-based solutions. We previously also discussed the technical benefits and disadvantages of the two technologies.

All in all, our conclusion is this: Use the right storage for the right job! Use disk and virtual tape libraries only where they are useful and have advantages over physical tape. Minimize your power and cooling costs if at all possible.

2.4 Summary

In this chapter, we focused on planning and sizing issues that you must take into consideration before implementing Tivoli Storage Manager.

3

Project planning

In this chapter, we discuss the necessary planning steps for IBM Tivoli Storage Manager V5.5 deployment scenarios. We differentiate the deployment into small, medium, or large environments. Our discussion of these environments provides a realistic understanding of different examples of Tivoli Storage Manager architecture designs.

When reading this chapter, remember that every deployment strategy is unique, and planning it properly is essential for a successful implementation.

We discuss the following topics:

3.1 Deployment environments

We cover four types of environments:

- Simple environment
  - Small installation
  - Maximum of 20 clients
  - Small database and mail servers
- Moderate environment
  - Medium-size installation
  - Maximum of 100 clients
  - Large database and mail servers
  - IBM Tivoli Storage Manager Storage Agent for LAN-free backups
- Complex environment
  - Large installation
  - Maximum of 500 clients
  - Integrated with IBM Tivoli Storage Manager for Advanced Copy Services
- Very complex environment
  - Huge installation
  - More than 500 clients
  - Multiple geographic sites

Note: Our classification here is based on the experience and efforts that are necessary to design and install an IBM Tivoli Storage Manager environment.

3.1.1 Simple environment

The managed system assets of a small installation typically are consolidated in a single location. Consider this example of a simple environment, with the following server types to back up:

- Microsoft Windows 2003 file server
- Windows 2000 and Windows 2003 application servers (DNS, DC, WINS, and so on)
- AIX V5.3 application servers
- Windows 2003 Exchange Server
- HP-UX V11.0 with Oracle databases

The installation consists of the following components:

- IBM Tivoli Storage Manager Server Basic or Extended Edition
- IBM Integrated Solutions Console (ISC) and Tivoli Storage Manager Administration Center
- IBM Tivoli Storage Manager Client
- IBM Tivoli Storage Manager for Mail (for Microsoft Exchange Server)
- IBM Tivoli Storage Manager for Databases (for Oracle databases)

Figure 3-1 provides an overview of a typical Tivoli Storage Manager solution for a small environment.

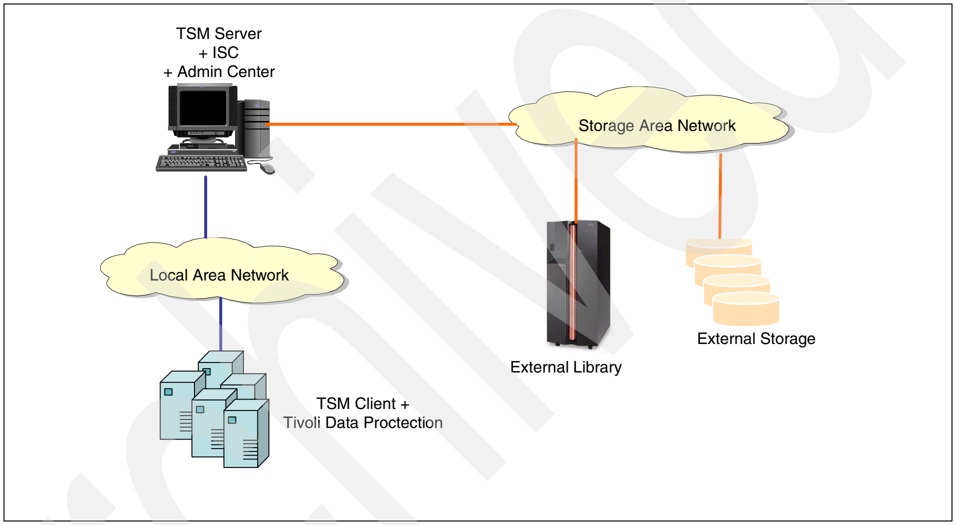

Figure 3-1 Simple environment

In this scenario, we recommend you install one Tivoli Storage Manager server and one database instance. On the same machine you can run both administrative consoles (ISC and Administration Center). It is better to use external storage due to performance issues and an external library to use copy storage pool functions.

This scenario typically does not include a configured disaster recovery plan, but it is important to notice that even in a small environment, the disaster recovery plan can help you recover your environment in an acceptable manner in case of a disaster. Depending on your hardware configuration, you can choose an environment that provides scalability.

Important: If you only have Microsoft Windows servers and applications such as Microsoft SQL Server and Microsoft Exchange, consider using IBM Tivoli Storage Manager Express. See:

http://www.ibm.com/software/tivoli/products/storage-mgr-express/

3.1.2 Moderate environment

The moderate environment is the most often used design for Tivoli Storage Manager solutions. In this scenario the hardware inventory is located in a single building and includes many servers and large databases.

This scenario typically uses Disaster Recovery Manager (DRM) along with tape vaulting to recover the environment from a major loss of IT infrastructure. For these environments with large databases, we recommend you have storage area network (SAN) connected tape library and devices to perform LAN or LAN-free backups. Using LAN-free backups, you can improve the overall backup performance and reduce the network traffic significantly.

This example scenario describes a moderate environment that requires the support of additional server types (the last three list items are the additional server types). The following servers must be backed up:

- Microsoft Windows 2003 file server
- Windows 2000 and Windows 2003 application servers (DNS, DC, WINS, and so on)
- AIX V5.3 application servers
- Windows 2003 Exchange Server
- HP-UX V11.0 with Oracle databases
- AIX 5.3 with SAP databases
- AIX 5.3 with Lotus Domino
- HP-UX V11.0 application servers

The installation consists of the following components, of which the last two are additional components that are not necessary in the simple environment example (see 3.1.1, “Simple environment” on page 88):

- IBM Tivoli Storage Manager Server Extended Edition
- ISC and Administration Center
- IBM Tivoli Storage Manager Client
- IBM Tivoli Storage Manager for Mail (for Exchange and Domino)

- IBM Tivoli Storage Manager for Databases (for Oracle databases)
- IBM Tivoli Storage Manager for ERP (for SAP databases)
- IBM Tivoli Continuous Data Protection for Files (CDP) (for Microsoft Windows file servers)
- IBM Tivoli Storage Agent for clients with large databases or mail servers

Figure 3-2 provides an overview of the Tivoli Storage Manager solution for a moderate environment.

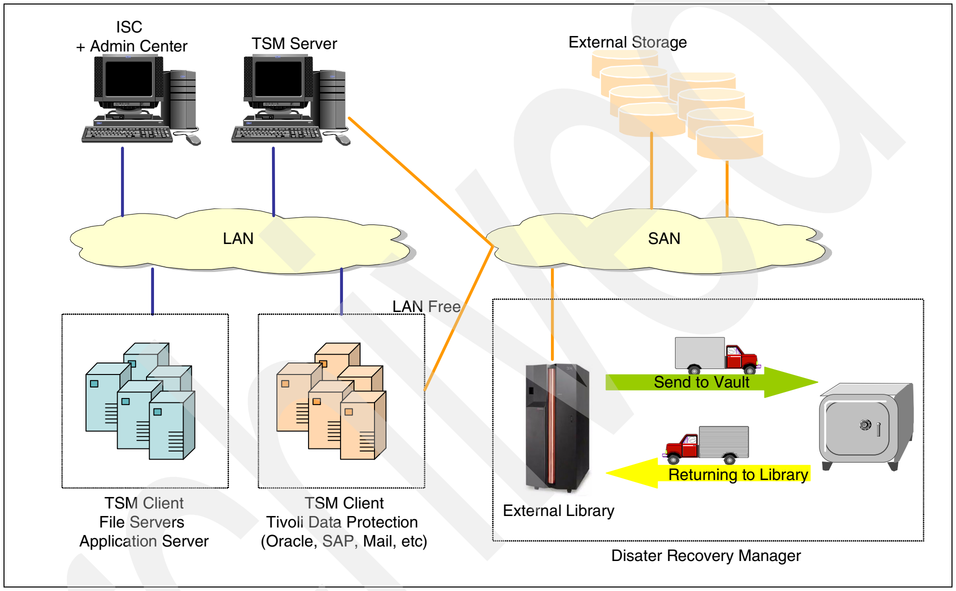

Figure 3-2 Moderate environment

In this scenario, we recommend you install the Tivoli Storage Manager server on one machine and the administrative console (ISC and Administration Center) on another machine. If your machine is powerful enough to support both applications, you can choose to have both on the same system. The external storage is necessary for the database, log, and storage pool volumes.

It is important to plan for Tivoli Storage Manager database growth according to the size of the environment. A database with more than 120 GB can cause performance problems for the environment, specifically backup performance, which is poorer. Consider using more than one instance if your Tivoli Storage Manager database grows to more than 80 GB. It is also necessary to estimate the external library, drives, and slots that permit the copy storage pool backups. You must set up those clients with large databases or mail servers to perform LAN-free backups.

3.1.3 Complex environment

Building on the fundamentals of a medium installation, a large installation focuses on scalability and on additional features for backing up some environments. The recommended hardware specification or higher is required to properly scale the infrastructure. The scope for a complex environment is almost the same as the scope for a moderate environment, but we can include integration with Copy Services and with FlashCopy on external storage to configure the backup, network-attached storage (NAS) server using Network Data Management Protocol (NDMP) or hierarchical storage management (HSM). For this type of environment, we still must use the DRM for tape vaulting. In addition, this environment usually includes more than one instance, and we must use the library manager and library client configuration.

> **Tip:** If you have an application that has intense I/O and CPU utilization and you want to reduce the utilization of these resources on production machines, we highly recommended that you use IBM Tivoli Data Protection for Advanced Copy Services during backup tasks to integrate the FlashCopy with Tivoli Storage Manager backup.

In this example scenario, which describes a complex environment, the following servers must be backed up (note that we list an additional five items):

- Microsoft Windows 2003 file server
- Windows 2000 and Windows 2003 application servers (DNS, DC, WINS, and so on)
- AIX V5.3 application servers
- Windows 2003 Exchange Server
- AIX V5.3 with SAP databases
- AIX V5.3 with Lotus Domino
- HP-UX V11.0 with Oracle databases
- HP-UX V11.0 application servers
- Network attached storage server
- Linux applications
- AIX V5.3 with inactive data that must be managed
- AIX V5.3 with SAP and DB2 integrating with FlashCopy backup in IBM DS8300 storage system
- Windows 2003 with SQL Server

The installation consists of the following components (we list an additional component last):

- Tivoli Storage Manager Server Extended Edition
- ISC and Administration Center
- Microsoft Windows 2003 with SQL Server
- IBM Tivoli Storage Manager Client
- Tivoli Storage Manager for Mail (for Exchange and Domino)
- Tivoli Storage Manager for Databases (for SQL and Oracle databases)
- IBM Tivoli Storage Manager for ERP (for SAP databases)
- Tivoli Continuous Data Protection for Files (CDP) (for Windows file servers)
- Storage Agent for clients with large databases or mail servers
- Tivoli Storage Manager for Advanced Copy Services (FlashCopy integration)

Figure 3-3 provides an overview of a Tivoli Storage Manager solution for a complex environment.

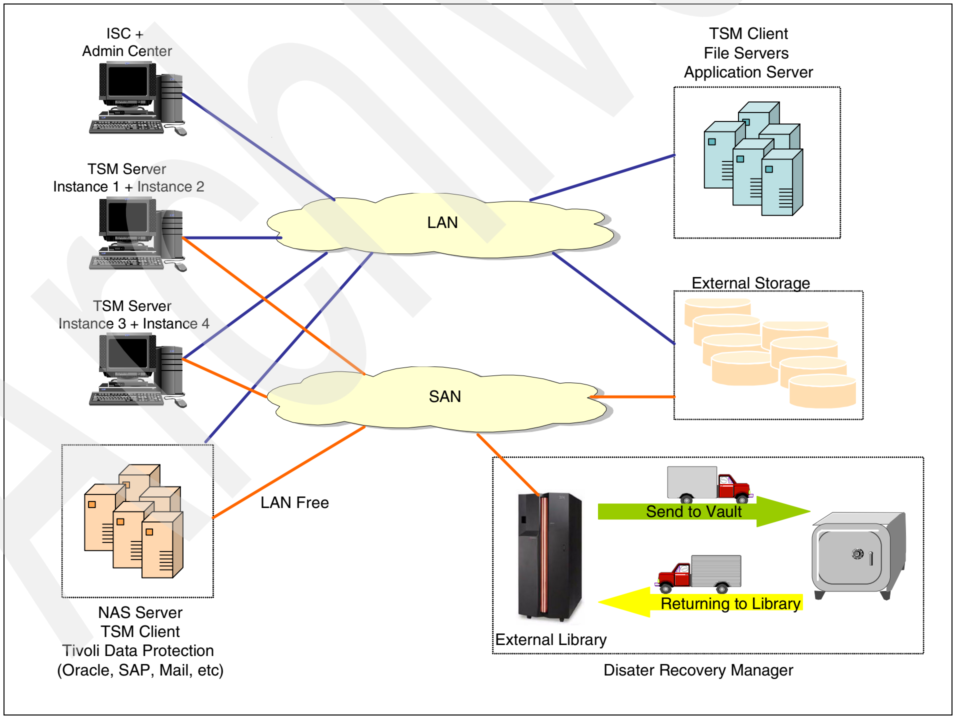

Figure 3-3 Complex environment

You must be aware of several issues concerning this scenario, so you must perform an accurate planning and assessment stage before implementing this large installation. It is critical that you map all components in the topology to the recommended hardware specifications in order to achieve a highly distributed environment with realistic goals. We recommend you thoroughly study and understand the client environment before proceeding and implementing any architectural design.

It is also important for you to account for all variables within the topology. Give substantial consideration to the infrastructure hardware requirements and the underlying network topology. You must assess network bandwidth, latency, and firewall restrictions. We recommend you install each Tivoli Storage Manager Server instance on a separate machine and the administrative console (ISC and Administration Center) on a different system. If your machine is powerful enough to support both applications, you can choose to have both on the same system. However, remember that when you use the same machine you share all resources, which might impact the performance. For performance reasons, external storage is mandatory for the Tivoli Storage Manager database, log, and storage pool volumes, for performance reasons. As we have already mentioned, if you plan to have DRM, you are required to perform capacity planning on all your storage devices, libraries, and drives. You must configure these clients with large databases or mail servers that do not use FlashCopy functions, using LAN-free backup.

3.1.4 Very complex environment

This type of scenario is also relatively common. It typically involves multiple sites with required hardware. This type of solution is very expensive but provides good protection for critical data. In this environment, you can have the same type of servers that we described earlier in the discussion of a complex environment, but for a complete disaster recovery solution, this scenario uses a server-to-server configuration and duplicates the backup in different sites. See Figure 3-4 on page 95.

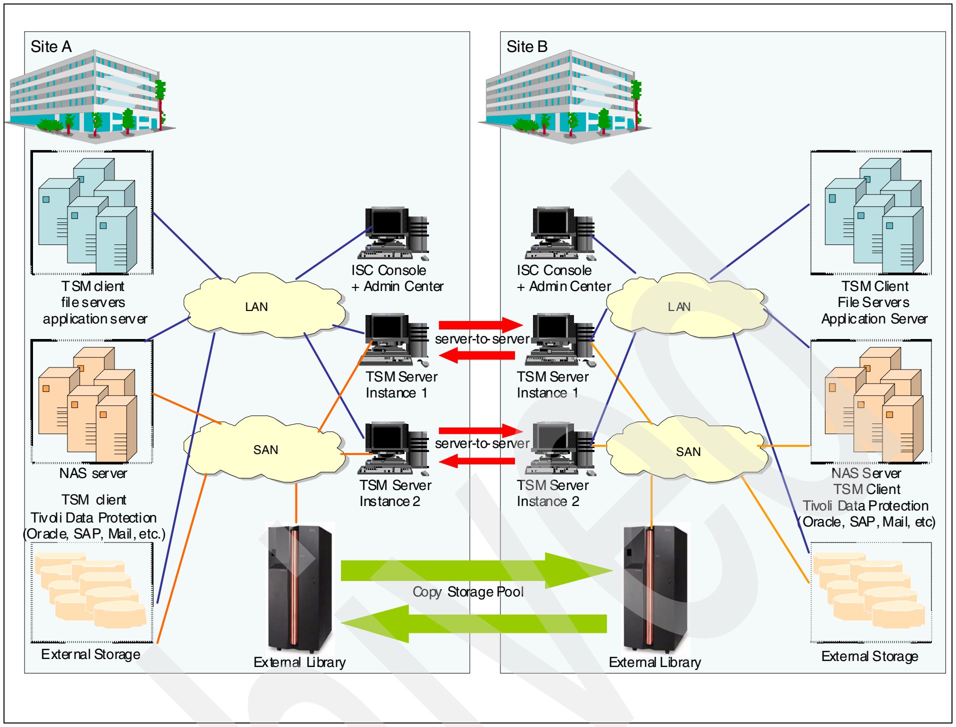

Figure 3-4 Very complex environment

This specific scenario presents almost the same issues as a complex environment. The difference between them is that we must consider more than one geographic site, so in most cases, the issues are related to hardware requirements. A network link between both sites is necessary, and depending on the distance between sites, you must consider utilizing a technology that supports this type of communication, for example, an extended fabric or Dense Wavelength Division Multiplexing (DWDM).

You can have more than one instance on each site, and only one console is necessary for all Tivoli Storage Manager servers. The external storage is necessary for both Tivoli Storage Manager servers. Also, this scenario requires capacity planning of all your storage devices, library, and drives.

3.2 Preparing for deployment

The deployment of any product is made easier when you take into consideration all relevant issues and devise a complete plan before rolling it out. It is important that you invest time and effort to reach an understanding of how Tivoli Storage Manager is going to be used and the possible scenarios you can create that best fit your environment.

The key to achieving a successful Tivoli Storage Manager deployment depends on adhering to a common road map, tailored to the client requirements.

The purpose of this section is to describe the groundwork and documentation required for a successful deployment.

3.2.1 Solution overview

The solution overview helps you understand the solution concepts and business value and the system architecture considerations. You must consider the following items:

- Solution concepts and business value
- Solution architecture
- Recommended software, hardware, and tools

3.2.2 Solution planning guide

The solution planning guide helps you to plan for the client engagement. You must consider the following items:

- Skills required to implement the solution
- Spreadsheet of time estimates to plan for a services engagement
- Task descriptions
- Best practices

3.2.3 Implementation guide

The implementation guide is a technical guideline that helps you set up the solution and learn all the important techniques required to implement a customized solution. You must consider the following items:

- Implementation checklist
- Instructions, tips, and installation best practices
- Customization information
- Use cases that can be used as part of a demonstration or education
- Sample code, data, scripts, and configurations

3.3 Identifying resources

To be able to successfully deploy a Tivoli Storage Manager solution requires considerable research and collaborative effort between technicians, managers, consultants, administrators, and users.

The purpose of this section is to identify the required skills and where you can find available resources to help with Tivoli Storage Manager deployment.

3.3.1 Implementation skills

You must establish a team of people with a variety of skills, depending on the solution that you are implementing. The following list describes the skills your team requires in order to design, develop, and deploy a Tivoli Storage Manager solution:

- Expert operating system administration skills

  Configure system user profiles and file security, tune the OS to improve the backup performance, and install Tivoli Storage Manager file sets based on the platform. For a full list of the supported platforms see:

  http://www-306.ibm.com/software/tivoli/products/storage-mgr/platforms.html

- Expert connectivity and networking skills

  Establish communication between all the machines involved in the solution.

- Basic storage area network (SAN) skills

  Configure the SAN switch to permit your servers to access the tape drives, external storage disks, and allocated devices.

- Basic external storage skills (disk subsystems)

  Format and allocate new volumes for the Tivoli Storage Manager solution. For example, if you are using external storage for database, log, and storage pools volumes, you require someone on the team who has the ability to implement and configure external storage such as IBM DS family, EMC, HDS, and so on. This requirement depends on your external storage system.

- Expert tape technologies and devices skills

  Ability to configure tape library and drives. Tivoli Storage Manager supports many different libraries and tapes. Access® the following links for detailed information:

  http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_AIXHPSUNWIN.html

  http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_iSeries.html

  http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_Linux.html

- Expert applications skills

  Ability to configure a backup for a specific application, whether or not you are using IBM Tivoli Storage Manager for Databases. For example, while backing up an Oracle database using Tivoli Data Protection for Databases, which includes Oracle support, it is necessary that the database administrator create the RMAN scripts to use in conjunction with the IBM Tivoli Storage Manager API (recovery manager).

- Expert Tivoli Storage Manager server and data protection skills (mail, database, and applications)

  Ability to plan, install, configure, and execute administrative tasks and perform problem determination on Tivoli Storage Manager and all installed optional components.

For example, if you intend to use internal disks only for database, log, and storage pool volumes, you do not require the skill to configure external disk storage in your solution.

Note: It is important to identify all the applications and resources that exist in the client environment prior to defining the skills necessary for your deployment plan.

3.3.2 Available resources

The prerequisite skills described in 3.1.1, “Simple environment” on page 88 are those required to customize and develop the solution. For each skill, a variety of resources are available to help acquire the necessary skill level or become more familiar with the products involved. We have the following Tivoli Storage Manager educational resources available:

- Online help

  Access IBM Tivoli Storage Manager online product manuals at:

  http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp

- Training

  IBM PartnerWorld® provides current information about available classes, their dates, locations, and registration. Additionally, check the PartnerEducation Web site:

  http://www.ibm.com/software/sw-training/

  Consider using the IBM Virtual Innovation Center™ at:

  http://www.ibm.com/partnerworld/vic

  You can access IBM Learning Services at:

  http://www.ibm.com/services/learning/index.html

- IBM Redbooks publications

  You can access practical and architectural information regarding IBM hardware and software platforms from the IBM Redbooks Web site, including information about Tivoli Storage Manager concepts and architecture at:

  http://www.redbooks.ibm.com

3.4 Solution scope and components

It is important to understand the scope of the solution. The solution can be a basic Tivoli Storage Manager server and client solution, or you can add additional components and make it more complex. In this section we provide tips to identify all the solution components. The discussion includes the following topics:

- 3.4.1, “Defining solution timing” on page 100
- 3.4.2, “Solution deployment tasks” on page 101

3.4.1 Defining solution timing

> **Important:** Whenever you estimate the time required to perform the various tasks associated with a solution deployment, it is important to document the assumptions related to the environment and the skills of the team involved.
>
> For instance, the time required for a Tivoli Storage Manager deployment can depend on:
>
> - Knowledge of IBM Tivoli Storage Manager Extended Edition terms
> - Knowledge of IBM Tivoli Storage Manager for Databases terms
> - Knowledge of data backup and recovery planning
> - Knowledge of systems, storage, and networks
> - Previous experience with products
> - Experience in prior deployments, and so on

The basics of deploying a profitable Tivoli Storage Manager solution is to correctly identify the tasks that you must perform and adequately allocate the necessary time to perform them. This section provides guidance about the tasks that you might require to perform in order to implement a Tivoli Storage Manager server and client solution, and we provide guidelines for timing estimates. The estimates rely largely on the following basic assumptions:

- Solution complexity

  The number of different Tivoli Storage Manager complementary products that are introduced into the solution (Tivoli Storage Manager for data protection, Tivoli Storage Manager for SAN, Tivoli Storage Manager for Space Management, and so on) relate directly to the complexity of the solution and to the time required for:

  – Designing the solution
  – Identifying performance bottlenecks
  – Implementing the solution

- Solution size

  The number of clients, Tivoli Storage Manager instances, and hardware components correspond to the time required for installing software and hardware.

- Service level agreement (SLA) requirements

  When the requirement for faster restores and recovery arises, products such as IBM Tivoli Storage Manager for Storage Area Network and IBM Tivoli Storage Manager for Hardware are implemented. This impacts heavily on the time required for designing and implementing the solution.

Use Table 3-1 to list and document sizing assumptions.

Table 3-1 Assumptions

Assumption	Values
Number of Tivoli Storage Manager server instances	
Number of backup/archive clients	
Number of Tivoli Storage Manager data protection clients	
Number of local area network or LAN-free installations	
Number of Tivoli Storage Manager for Advanced Copy Services	
Complexity of disaster recovery scenario	
Degree of implementation of SAN-attached disk and tape drives	

3.4.2 Solution deployment tasks

This section discusses the tasks necessary for solution deployment. You can use Table 3-2 to document the timing estimates for the major deployment tasks. These tasks are broken down into subtasks in later sections and discussed in further detail.

Table 3-2 Major solution deployment tasks

Task	Estimated time (hours)
Plan the solution	
Implement the solution	
Test the solution	
Establish performance and capacity baseline and reporting	
Client training	
Close the engagement	
Total	

Plan the solution

When planning the deployment of the Tivoli Storage Manager solution, you must include the subtasks shown in Table 3-3.

Table 3-3 Solution planning subtasks

Task	**Estimated time (hours)**
Gather detailed information about client environment	
Analyze client requirements	
Design the solution	
Develop implementation plan	
Develop a test plan	
Develop a plan for reporting	
Analyze client education requirements	

The tasks listed in Table 3-3 are described in more detail in the following sections.

Gather detailed information about client requirements

At the beginning of your engagement, you meet with your clients to understand their proposed objectives and gather their requirements. At first, you determine the functional requirements. Functional requirements define the business functions that the Tivoli Storage Manager system is going to provide. You determine the requirements by developing a good understanding of the business requirements and of what you hope to achieve. It is important to gather these requirements early in the process and discover any challenges that might lie ahead while you can still deal with them easily. Once you have determined the functional requirements, you can clarify the technical and system requirements.

Obtaining technical requirements involves spending time at the client site to determine and understand the client's hardware and software installation and infrastructure. To obtain the technical requirements, complete the following tasks:

1. Complete a high-level inventory of the environment.
2. Inventory everything that the solution might be called on to recover (data, servers, disks, databases, and applications).
3. Note software levels for OSs, databases, and applications.
4. Inventory every source that the solution might depend on (network segments, disk arrays, tape libraries, and vaulting services).

The system requirements are also determined using the inventory. Complete the following steps:

1. Note the sizes of file systems, disks, and databases, as well as the number of files on client servers.
2. Evaluate expected growth and nightly changes.
3. Review the client's data retention obligations (such as the Sarbanes-Oxley regulations).
4. Conduct additional inquiries to obtain information about backup and restore windows for each client system because this information has a direct effect on the SLA with the client. This information also dictates the size and performance requirements of the Tivoli Storage Manager solution.

Analyze client requirements

Compile and review the client requirements information you have gathered. Create the appropriate diagrams and charts in order to study the flow of data during backups and recoveries.

Design the solution

After the environment has been inventoried and described in sufficient detail, and all requirements have been agreed to and documented, you can complete the detailed design of the backup and recovery solution. The solution design phase focuses on data flows and on an analysis of how copies of data on disks are transferred through various networks, or SAN, to sequential disk or tape for final storage. The solution design phase focuses on operational schedules, backup schedules, and an analysis of how various operations and various data transfers can complete in the allotted time slots for backup and recovery.

In designing the solution, you choose the appropriate hardware to support the requirements along with the right type and number of Tivoli Storage Manager client products. It is important to design the solution to support the most demanding component. This component may be a large file server with millions of files, a database server with many terabytes of data, tight recovery windows that demand the requirement for LAN-free backup, or the use of Tivoli Storage Manager for Advanced Copy Services.

For example, disaster recovery requirements are a significant concern during the design phase. No matter how complex, an entire backup and recovery infrastructure is but one component of the larger disaster recovery solution. If the disaster recovery requirements call for off-site recovery from site disasters, you must design and size the backup and recovery solution to support such recoveries. When you have identified and taken into account the few most demanding components, designing the rest of the solution around these components is straightforward.

The solution design might require revisiting the requirements analysis phase due to unfavorable costs that arise while trying to satisfy the client's initial recovery or data retention requirements. In such a situation, you can lower the requirements to balance requirements and costs. For further discussion on solution design, see 2.2, "Solution architecture considerations" on page 43.

Develop implementation plan

Successfully implementing a backup and recovery infrastructure requires a fair amount of coordination and collaboration. We recommend undertaking at least a small amount of project planning and project management in order to implement a detailed solution with minimal impact on regular production operations. As early in the process as possible, develop plans for procuring the actual hardware, software, services, and supplies. Engage technical specialists to take the lead in integrating the Tivoli Storage Manager server and the Tivoli Storage Manager clients with the LAN, the SAN, and the various client servers. For complex implementations, you must use an experienced Tivoli Storage Manager solution architect. For further discussion on acquiring the right skills and resources, see 3.3, "Identifying resources" on page 97.

Develop test plan

Develop a comprehensive test plan to test the Tivoli Storage Manager implementation. The test plan includes testing in a production or production-like environment to minimize unforeseen bottlenecks and impacts when the new backup and restore solution is put into production. You must consider data relationships between different systems or applications. Define measurable criteria for successful testing, and enter the criteria into the test plan. The plan includes testing solution infrastructure (correct data flow through network and SAN), testing the Tivoli Storage Manager server configuration (data retention in accordance with SLA, storage hierarchy, clients, and administrative schedules), and testing disaster recovery. The plan also includes testing basic client backup and restore functionality.

Develop plan for reporting

Reporting is an important component of successfully running a Tivoli Storage Manager solution. We recommend doing daily and monthly reporting, performing trend analysis, and creating an initial baseline for reporting. Create a plan for testing the baseline of the Tivoli Storage Manager server, the server capacity, and of client performance. The outcome of the test helps ensure that the daily support staff has a baseline to work from when investigating performance or capacity issues in the Tivoli Storage Manager infrastructure (Tivoli Storage Manager server, clients, network, and SAN).

Collect the same performance and capacity data on a monthly basis, analyze trends, review the data, and act upon it if required. Create a plan for reporting on the daily and monthly status of client backups (successes and failures). The Tivoli Storage Manager Operational Reporting feature automates some of the monitoring tasks you typically perform on your clients and server. Operational Reporting is a free tool that runs on Microsoft Windows platforms and can report on any type of Tivoli Storage Manager server. However, it is not a trending tool.

For trend analysis, we recommend using the Open Database Connectivity (ODBC) application programming interface (API) that is included in the Tivoli Storage Manager Client installation package. You can use this API to access data directly in the Tivoli Storage Manager server database using SQL SELECT statements. After you install the API driver, you can use a spreadsheet or database application that complies with ODBC to access the Tivoli Storage Manager server database for information.

Analyze client education requirements

The long-term success of any backup and recovery solution requires that the support staff is capable of managing the solution effectively. Daily backup operations are sensitive to changes in the production environment. The support staff must be well trained in the details of Tivoli Storage Manager administration. You must create client education plans.

Implement the solution

You can implement the solution using the tasks described in Table 3-4.

Table 3-4 Solution implementation subtasks

Task	Estimated time (hours)
Hardware installation	
Tivoli Storage Manager server installation and configuration	
ISC and Administration Center installation	
Client installations and configurations: ► Basic backup/archive clients ► LAN-free clients ► Data Protection client modules ► Tivoli Storage Manager for Advanced Copy Services	

Hardware installation

You install and configure the hardware included in the solution. This hardware includes the physical Tivoli Storage Manager server, network, SAN, disk subsystems, libraries, and tape drives.

IBM Tivoli Storage Manager server installation and configuration

Install and configure Tivoli Storage Manager V5.5. Configuration includes creating storage hierarchies, policy domains, and housekeeping schedules. See Chapter 4, "Installation and configuration" on page 119 for further details on installing Tivoli Storage Manager V5.5 in a Microsoft Windows or a UNIX environment.

Integrated Solutions Console and Administration Center installation

You can use the Integrated Solutions Console (ISC) and IBM Tivoli Storage Manager Administration Center to administer Tivoli Storage Manager V5.5. We recommend using these tools because they provide centralized administration of all your Tivoli Storage Manager instances and a comprehensive view of the general health of your Tivoli Storage Manager servers. See 4.6, "Integrated Solutions Console" on page 187 for further details on installing ISC and the Administration Center on Microsoft Windows.

Client installation and configuration

After you have implemented the Tivoli Storage Manager server and its supporting network and storage infrastructure, the project proceeds to the more delicate work of integrating all data and applications into the new backup operations. Exercise care when configuring clients because correct configuration, in accordance with SLAs and disaster recovery plans, is the only thing that stands between a valid and successful backup and a backup that cannot be used to recover critical data or applications. Where you encounter traditional configurations (backup and restore clients on simple systems), implementation proceeds directly and carefully. Where you encounter unique or complex configurations (backup and restore clients on file servers or complex application servers, database servers, mail servers, or LAN-free clients), you must perform extra testing. Where you encounter large numbers of similar systems, you might test a standard client installation and replicate or push it to all clients using application deployment software, such as Tivoli Configuration Manager.

Test the solution

The tasks required to test the solution are listed in Table 3-5.

Table 3-5 Solution testing subtasks

Task	Estimated time (hours)
Perform backup and restore test plan	
Perform disaster recovery test plan	

Perform backup and restore test plan

Test the solution infrastructure, Tivoli Storage Manager server configuration, and basic client backup and restore, and note the results for later documentation.

Perform disaster recovery test plan

An important part of the entire deployment is testing to validate that the new solution can successfully recover the entire client environment in case of total site failure. A complete recovery is more complex than recovering single systems at a time. Clients with large server environments often use very complex applications that run on and make use of data from many servers at the same time. With such applications and servers, it might be important to recover in the right order to avoid inconsistent data or loss of data. Complete the disaster recovery test plan and note the results for later documentation.

Establish performance and capacity baseline and initiate reporting

Tasks necessary for establishing a baseline regarding performance and for initiating reporting on the Tivoli Storage Manager system are listed in Table 3-6.

Table 3-6 Solution baseline and reporting subtasks

Task	Estimated time (hours)
Gather baseline material regarding Tivoli Storage Manager server, server capacity, and client performance	
Implement reporting	

Gather baseline material regarding performance and capacity

Carry out various performance tests, such as testing backup and restore speed across network and SAN, and Tivoli Storage Manager server data movement operations. Note for later documentation the following capacity parameters:

- Allocated disk and tape cartridges
- Tape drive utilization
- Tivoli Storage Manager server database and log utilization
- Tivoli Storage Manager Client initially stored data and nightly backup volume

Implement reporting

Implement reporting using the agreed-upon parameters and reporting tools.

Client training

The tasks required to perform client training are listed in Table 3-7 on page 108.

Table 3-7 Client training subtasks

Task	Estimated time (hours)
Complete client training	

We recommend that key system administrators be formally trained in the use of Tivoli Storage Manager. Other system administrators and database administrators can participate closely in the technical configuration of the solution and receive hands-on training from the Tivoli Storage Manager Technical Specialists who are implementing the solution.

IBM offers extensive courses in Tivoli Storage Manager. We recommend that technical staff responsible for administrating the environment on a daily basis attend the two basic courses “Tivoli Storage Manager V5.5 Implementation and Administration Part One” and “Tivoli Storage Manager V5.5 Implementation and Administration Part Two.” Senior technical staff might also attend “Tivoli Storage Manager V5.5 Advanced Administration, Tuning and Troubleshooting” to enhance their troubleshooting skills.

Workshops are valuable alternatives to official courses. The available dates for the courses might not fit into the client’s schedule, and Tivoli Storage Manager Technical Specialists can successfully arrange and run a customized workshop. A customized workshop also provides the opportunity to tailor the training to the client’s specific needs.

Close the engagement

Closing the engagement is performed using the tasks described in Table 3-8.

Table 3-8 Engagement closing subtasks

Task	Estimated time (hours)
Document solution	
Close engagement	

Document solution

A Tivoli Storage Manager solution is a complex implementation that you must fully document. Document all configuration (software and hardware) parameters and daily production procedures. Gather and document output from the carried out backup and restore test plan and disaster recovery test plan along with documenting information about the disaster recovery plan itself. Provide reporting baselines and reporting procedures.

Final verification

Conduct a final verification to ensure that the implemented solution matches the initial client requirements. Hand over the documentation to the appropriate parties.

3.5 Strategies for a successful implementation

Building and sustaining a storage management solution can challenge even the most successful organization. As mentioned earlier, make sure your deployment team includes the right personnel, carefully plan your deployment, and receive commitment from the organization as a whole. In addition to these strategies, we discuss other key considerations that can help you achieve a successful implementation of the Tivoli Storage Manager in this section.

3.5.1 Read updated documentation

It is important to always verify the online information about Tivoli Storage Manager because it is updated frequently. The readme file, which is shipped with the product, contains the release notes and features that come with the latest version.

3.5.2 Obtain commitment

Experts agree that a storage management initiative stands little chance of success unless upper-level management takes ownership of the project. Project leaders can initiate the effort, but commitment from upper-level management to get everyone to support it is an essential ingredient.

Note: The deployment team must ensure that mechanisms are put in place to facilitate user consultation and feedbacks.

3.5.3 Verify technology integration issues

As with any complex system, Tivoli Storage Manager can touch many aspects of your computing infrastructure. It is important to plan for these technology integration and interoperability issues early on to ensure a successful roll-out. The infrastructure design and network directly impacts the Tivoli Storage Manager solution, so all those aspects must be carefully examined and considered during your planning stage.

3.5.4 Clearly establish the scope of the project

Establish the scope of any knowledge management project early on, try to be realistic in what you can achieve, and ensure that it truly provides business value. Companies sometimes deploy knowledge management systems hoping that they might eventually discover a purpose for the system. You must deploy aTivoli Storage Manager solution, like any knowledge management effort, to address a clearly defined business problem and to produce clearly defined benefits to your organization.

3.5.5 Educate the client

It is widely accepted that a favorable corporate culture is a prerequisite for the successful implementation of knowledge management initiatives, so you have to focus not just on technology, but also on the client culture and process. Sustainable knowledge management requires the introduction of a knowledge-sharing culture organized around key skills and a set of incentives to reuse knowledge to reinforce the skills required for the effective operation of the organization. You might not agree entirely with this point of view, but being aware of what is being written about knowledge management and considering how knowledge sharing and information reuse is currently practiced in your organization might reveal the type of culture you have to work with when you draw up your knowledge management program. A key success factor in the implementation of Tivoli Storage Manager is encouraging the right mind-set in the staff that is responsible for delivering a knowledge management solution to the organization.

3.5.6 Involve, train, and inform

The benefits of deploying a knowledge management solution are likely to be lost on most users in an organization because you cannot point to any immediate result. Tivoli Storage Manager requires a certain amount of education, user involvement in pilots, training, and promotion to ensure acceptance and effective use of the product.

In preparing your deployment plan, allocate time for the following tasks:

- Technical design of the solution
- Process of learning how to install, configure, and maintain the software
- Problem determination process
- Process of learning how to understand the application log
- Process of learning how to develop the fault diagnosis checklist
- Process of learning how to use the solution

Provide a skill-evaluation checklist so you can develop education modules according to the skills of client employees.

It is better to provide the client with some hands-on training instead of just showing a presentation. If training time is limited, you need to show the client the real application environment so that the client can learn the system sufficiently. To present the live demonstration, you must prepare some data for the application in advance and design an exercise that illustrates the use case steps. Users are understandably apprehensive about the generation of knowledge and information that can be associated back to individuals.

3.5.7 Obtain the right hardware

You might experience numerous problems if inadequate hardware resources have been allocated. Tivoli Storage Manager has significant resource demands, particularly in a large distributed implementation. You must resist using second-hand hardware that does not meet the minimum system requirements.

3.6 Sample case study

In this section we describe a sample case study for planning the time required for deploying Tivoli Storage Manager in a complex environment. We use the tables from 3.4, “Solution scope and components” on page 99.

It must be noted that the durations of task times provided in this section are examples and might vary greatly based on the actual environment and the skills of the team involved.

Our sample case study is of a company that offers a complete telecommunication solution that operates in a single main data center and has more than fifteen local sites. This company has many servers with large databases on various platforms, such as AIX, Solaris, and Microsoft Windows servers. Many UNIX and Windows servers work as applications servers. Calendaring and e-mail services are provided by the Exchange Servers. The database, file, and mail servers are connected through a SAN and share four IBM DS8100. The smaller servers use internal disks.

Currently, the backup solution shows some performance problems, and the client is not comfortable with it. Many different products provide installed platform- and application-specific backup capabilities. Administration of this type of solution is difficult: The backup window often expires, the restore is not reliable, disaster recovery tasks are manual, and storage media, especially tapes, are not managed efficiently.

First we must perform an environment inventory. After collecting all information about the servers, applications, external storage, tapes, SAN, and LAN, we can define the assumptions in Table 3-9 based on our sample case study.

Table 3-9 Defining assumptions

Assumption	Values
Number of Tivoli Storage Manager server instances	4
Number of backup/archive clients	210
Number of Tivoli Storage Manager Data Protection clients	28
Number of LAN-free installations	23
Number of Tivoli Storage Manager hardware installations	5
Complexity of disaster recovery scenario	High
Degree of implementation of SAN-attached disk and tape drives	High

Due to the large number of clients and different backup policies, our solution consists of four Tivoli Storage Manager instances and uses the library manager and library client configuration to share the library. All clients must have the Tivoli Storage Manager Client installed. For applications servers and small servers, the backup goes through the LAN. For large databases that are highly I/O and CPU intensive, and where performance is an issue, we address this issue using FlashCopy integrated with Tivoli Storage Manager. For the other large databases, we implement a LAN-free backup solution with the respective Tivoli Storage Manager for Databases according to the application.

After defining all the assumptions, we can now determine the efforts for all the consolidated tasks that we perform in the solution. Table 3-10 represents the total time that we plan to spend in order to implement the solution, based on our example case study. Given the statistics in Table 3-9, we have a better idea about the estimated time for solution deployment in a specific environment. See the tasks described in Table 3-10.

Table 3-10 Solution deployment tasks

Task	Estimated time (hours)
Plan the solution	196
Implement the solution	465
Test the solution	240
Establish performance and capacity baseline reporting	104

Task	Estimated time (hours)
Train the client	80
Close the engagement	44

Let us dig a little deeper into these tasks and their subtasks. As we have emphasized earlier, the planning phase is important because it helps define the scope of the deployment. See the subtasks described in Table 3-11.

Table 3-11   Plan the solution

Task	Estimated time (hours)
Gather detailed information about the client environment	20
Analyze client requirements	16
Design the solution	100
Develop an implementation plan	16
Develop a test plan	24
Develop a plan for reporting	16
Analyze client education requirements	4

When implementing the solution, many activities can be executed in parallel. So at certain points during the implementation, you can save time in executing these activities. See the tasks described in Table 3-12.

Table 3-12   Implementing the solution

Task	Estimated time (hours)
Hardware installation (server, library, SAN, external disks)	32
Tivoli Storage Manager server installation and configuration	18
ISC and Administration Center installation	6
Clients installation and configuration (one hour for each client)	210
LAN-free configuration (three hours for each client)	69
Tivoli Storage Manager for Databases configuration (two hours for each client)	56

Task	Estimated time (hours)
IBM Tivoli Storage Manager for Hardware (ten hours for each client)	50
DRM implementation	24

The plan tests must be performed for all files and folders or the whole database and applications for each client; all backup and restore tests must also be performed for some files and folders and for a few table spaces, just to validate the communication, performance, and configuration. See the tasks described in Table 3-13.

Table 3-13 Test solution

Task	Estimated time (hours)
Perform backup and restore test plan	40
Perform backup and restore test plan for a client with Tivoli Storage Manager for Databases	200

We estimate it takes 80 hours to analyze the collected data for capacity planning analysis. The implement reporting task helps collect data for administrative and future capacity analysis. See the tasks described in Table 3-14.

Table 3-14 Establish performance, capacity baseline, and reporting

Task	Estimated time (hours)
Gather baseline material regarding Tivoli Storage Manager server and client performance and Tivoli Storage Manager server capacity	80
Implement reporting	24

You might decide to provide formal training in a classroom or hands-on training. Sometimes it is important to help the Tivoli Storage Manager administrator to execute administrative tasks using best practices and all Tivoli Storage Manager capabilities. For this example we are using hands-on training and classroom training and allotting 40 hours for each. See the task described in Table 3-15.

Table 3-15 Client training

Task	Estimated time (hours)
Complete client training	40

To document the full solution and its implementation in this scenario, we spend 40 hours. This documentation includes all installation, configuration, and backup policies. To close the engagement, it is necessary for us to compile a final checklist to ensure all tasks have been performed and expectations have been met. See the tasks described in Table 3-16.

Table 3-16 Close the engagement

Task	Estimated time (hours)
Document solution	40
Close engagement	4

3.7 Summary

This chapter has covered key aspects of a successful deployment. For anyone who has deployed new technology to an enterprise, the concepts of careful planning and phasing a deployment through test and pilot environments are probably not new. Much of the content of this chapter might therefore be a quick refresher.

However, when knowledge management technologies are involved, it is important to always consider the cultural and organizational impacts. Finding a clear and defined purpose for your deployment, and then ensuring that the organization truly understands the requirement to solve that problem, is crucial for ensuring the organizational commitment to any required cultural changes.

Part 2

Deployment

In this part, we provide information regarding the deployment of IBM Tivoli Storage Manager V5.4. This part of the book includes the following chapters:

4

Installation and configuration

In this chapter we describe installation instructions for installing the IBM Tivoli Storage Manager and its related components in a small-to-medium environment. The instructions given in this chapter are not as detailed as the documentation in the Tivoli Storage Manager installation guides. However, the information provided here is the basic information you need to install Tivoli Storage Manager in Microsoft Windows, AIX, and Linux environments. We also cover Integrated Solutions Console (ISC) and Administration Center installation.

4.1 Small and medium environments

Installing IBM Tivoli Storage Manager includes the following activities:

- Installing the Tivoli Storage Manager server and related components on the server machine. You can choose to install the Integrated Solutions Console (ISC) and Administration Center on a separate machine or on the same machine where the Tivoli Storage Manager Server resides.
- Installing and configuring Tivoli Storage Manager client code on every machine that transfers data to Tivoli Storage Manager server-managed storage.

Before you install Tivoli Storage Manager, ensure that the prerequisites mentioned in 2.1, “IBM Tivoli Storage Manager system requirements” on page 40 are properly met. In addition, you must be familiar with the following operating systems, devices, and protocols:

- Operating system on the Tivoli Storage Manager server machine
- Operating system for Tivoli Storage Manager client machines
- Storage devices available to Tivoli Storage Manager
- Communication protocols installed on your client and server machines
- Any special system configurations you plan to use depending on your setup, such as Microsoft Active Directory®, Microsoft Cluster Server (MSCS), Veritas Cluster, and HACMP.

4.1.1 Installing in small and medium environments

This section describes how you can use this chapter for guiding the installation process of the Tivoli Storage Manager server and its related components. These instructions are not the only way to install Tivoli Storage Manager and related prerequisite components. You can follow these instructions to successfully install and set up a Tivoli Storage Manager environment for both Microsoft Windows and selected UNIX platforms:

1. Refer to the installation and configuration information for Microsoft Windows, AIX, and Linux provided in the following sections:
   - 4.2, “Microsoft Windows environment” on page 122.
   - 4.3, “IBM Tivoli Storage Manager server for AIX” on page 161.
   - 4.4, “IBM Tivoli Storage Manager for Linux Server” on page 174.
2. Once you have completed the installation, consult the administration tasks discussed in 4.5, “Getting started with administrative tasks” on page 181.

3. Refer to the additional instructions for installing the Integrated Solutions Console (ISC) and Administration Center in 4.6, “Integrated Solutions Console” on page 187.

4.1.2 Installing IBM Tivoli Storage Manager in sequence

Install the Tivoli Storage Manager server and related components on the server machine and Tivoli Storage Manager client code on every machine that you plan to use to transfer data to Tivoli Storage Manager server-managed storage.

We always recommend you install Tivoli Storage Manager in the following order:

1. Tivoli Storage Manager server component
2. Tivoli Storage Manager licenses
3. Tivoli Storage Manager language pack (if required)
4. Tivoli Storage Manager device drivers or the IBM device driver (for IBM Magstar® and Ultrium devices)
5. Integrated Solutions Console and the Administration Center
6. Tivoli Storage Manager backup/archive client
7. Tivoli Storage Manager operational reporting

We recommended that you install the administrative client on the local client where the Tivoli Storage Manager server is installed for administrative purposes. You also can install the administrative client on any client machine necessary to perform this function.

Apart from these components, you can also install the optional Tivoli Storage Manager Open Database Connectivity (ODBC) driver, which enables reporting applications to export data from the Tivoli Storage Manager database and customize the reporting result for your requirements.We do not cover these reporting products in this book.

Note: You cannot install both the Tivoli Storage Manager server component and storage agent component on the same machine.

4.2 Microsoft Windows environment

The following instructions provide an overview of how to install Tivoli Storage Manager in a Windows environment. For detailed instructions on how to install the Tivoli Storage Manager server in a Windows environment, refer to the *IBM Tivoli Storage Manager for Windows Installation Guide*, GC23-5973. This document can be downloaded from the publication Web site:

http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/topic/com.ibm.itsmfdt.doc/b_install_guide_windows.htm

The sections are:

4.2.1 Installing the IBM Tivoli Storage Manager server

To install the Tivoli Storage Manager server, follow these steps:

1. Insert the IBM Tivoli Storage Manager Windows Server® CD or start the self-extracting executable downloaded from the IBM Passport Advantage® Web site. Click **Tivoli Storage Manager Server**.

   The Choose Setup Language dialog is displayed.

2. Accept the default or select a language for the Tivoli Storage Manager installation dialogs and click **OK**.

   > **Note:** This selection only specifies a language for the Tivoli Storage Manager installation dialogs. You can specify full product language support by selecting the **Custom installation** option (mentioned later in this installation sequence) in the Tivoli Storage Manager Server InstallShield Wizard.

   The Tivoli Storage Manager Server installation window is displayed.

3. Click **Install Products**.
4. The Install Products dialog is displayed, as shown in Figure 4-1. Click **TSM Server**.

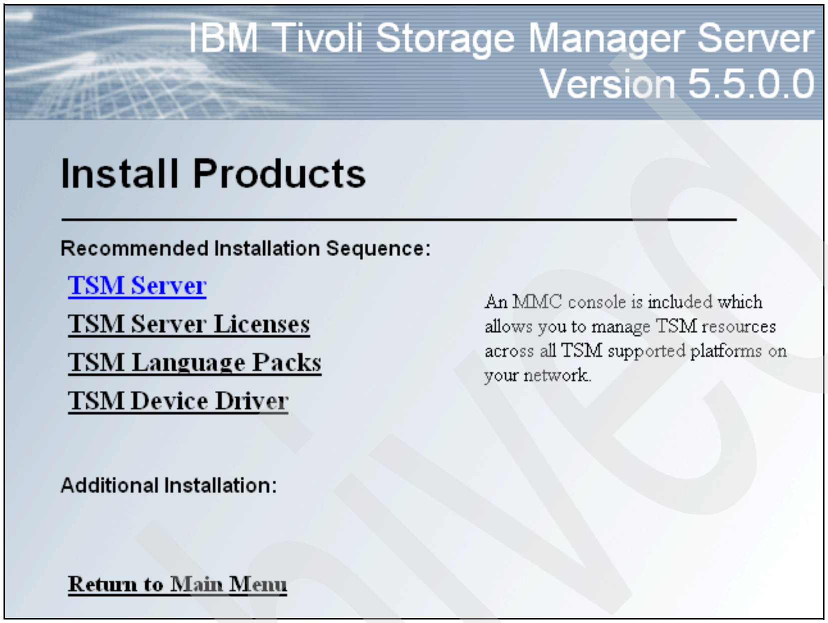

Figure 4-1 Install Products menu

5. The InstallShield Wizard window is displayed; click **Next**.
6. The License Agreement window is displayed. Read the license agreement, and accept the agreement. Click **Next**.

7. The Client Information window is displayed, as shown in Figure 4-2. Enter the user and organization information, and select one of the options under Install this application for. Click **Next**.

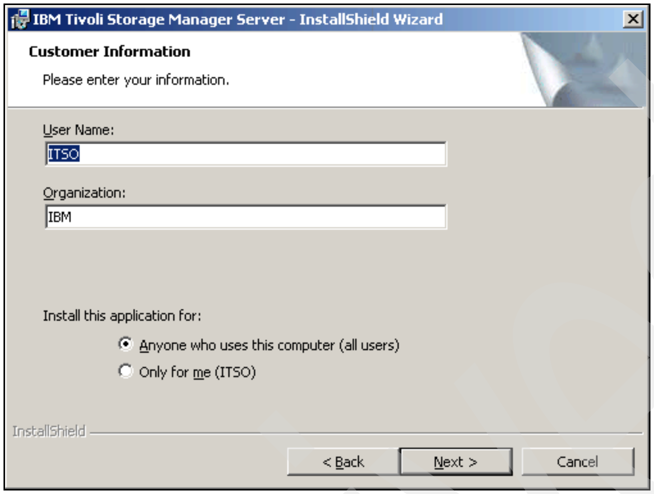

Figure 4-2 Client Information window

8. The Setup Type dialog box is displayed. Select **Complete** or **Custom**, and click **Next**.
9. Choosing Custom offers options such as selecting product language support, deselecting other Tivoli Storage Manager components, and choosing the installation path. When you have made your choices, click **Next**.
10. The Ready to Install the Program window is displayed. Click **Install**. The installation progress indicator is displayed on the next window.

11. After the installation completes, the InstallShield Wizard Completed dialog is displayed, as shown in Figure 4-3. Click **Finish**.

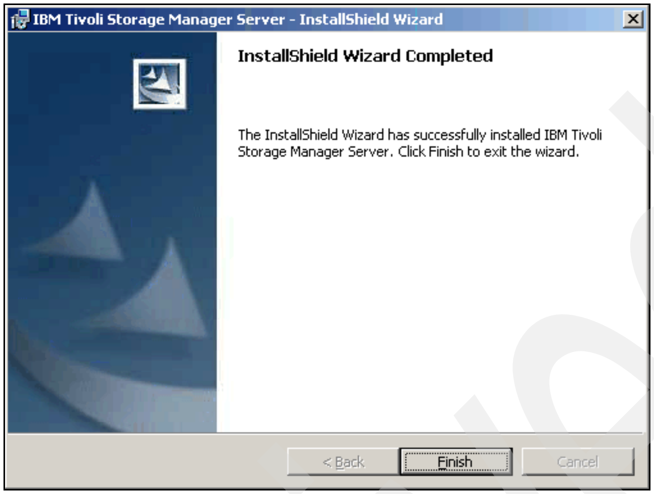

Figure 4-3 Installation complete

A dialog prompting you to restart your computer might be displayed. If you are not prompted, you are not required to restart.

4.2.2 Installing IBM Tivoli Storage Manager licenses

To install Tivoli Storage Manager licenses, follow these steps:

1. When you complete the server installation, return to the Install Products window (Figure 4-1 on page 123), and select **TSM Server Licenses**.
2. The InstallShield Wizard is displayed. Click **Next**.
3. The Client Information window is displayed. Enter the client name and the organization for whom the license was purchased, and select one of the buttons under Install this application for. Click **Next**.
4. The Setup Type dialog box is displayed. Select either **Complete** or **Custom** setup type.

Note: If you want to change the default installation directory, choose **Custom**. Otherwise the InstallShield Wizard chooses the default location for installing the files, which is normally your system drive.

5. The Ready to Install window is displayed. Click **Install**.
6. After the installation is finished, the InstallShield Wizard Completed dialog is displayed. Click **Finish**. You are not required to restart unless a restart is explicitly requested.

4.2.3 Installing IBM Tivoli Storage Manager device drivers

To install the Tivoli Storage Manager device drivers, follow these steps:

1. On the Tivoli Storage Manager Install Products window (Figure 4-1 on page 123), select **TSM Device Driver**.
2. Click **Tivoli Manager Device Drivers**.
3. The IBM Tivoli Storage Manager Device Driver InstallShield Wizard is displayed, as shown in Figure 4-4. Click **Next**.

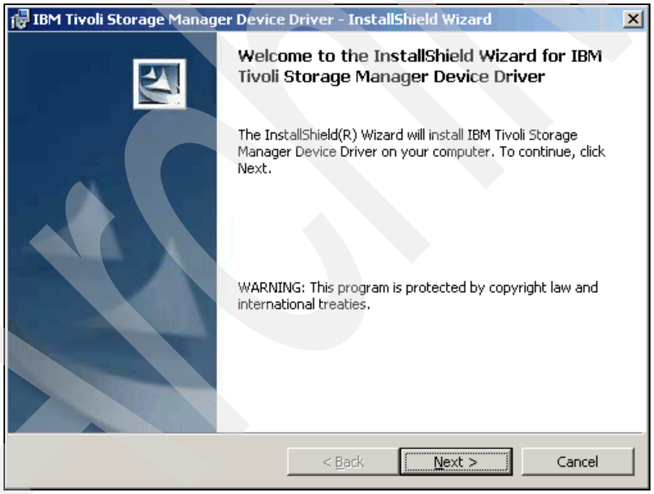

Figure 4-4 Device driver installation wizard

4. The Client Information window is displayed. Enter the user and client information in the respective boxes, and select one of the options under Install this application for. Click **Next**.
5. The Ready to Install window is displayed. Click **Install** to start the installation.

6. A progress indicator is shown, and when the installation completes, the Completing the Device Driver Installation Wizard window is displayed, as shown in Figure 4-5. Click **Finish**.

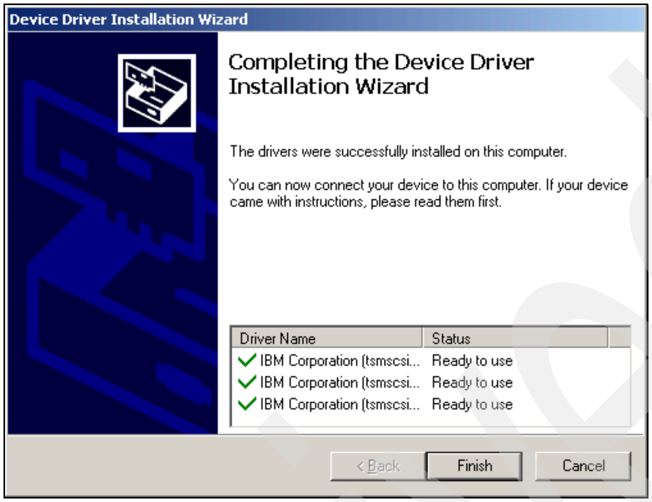

Figure 4-5   Device driver installation complete

7. An InstallShield Wizard Completed dialog box is displayed. Click **Finish**.

You have installed the Tivoli Storage Manager device driver, which is necessary to control IBM and non-IBM devices. If your environment consists basically of IBM devices, you must install the IBM Ultrium Device Driver. For more information see *IBM Ultrium Device Drivers Installation and User's Guide*, GA32-0430.

> **Tip:** The Removable Storage service in Microsoft Windows takes control of all tape and library devices. It is normally used by the NTBackup software. You must disable this service in order to avoid library initialization problems.

4.2.4 Configuring IBM Tivoli Storage Manager for a Windows server

The IBM Tivoli Storage Manager for Windows server provides a graphical interface called Tivoli Storage Manager Management Console tsmw2k, which includes a set of wizards that help you configure and manage your Tivoli Storage Manager system. One or more of these wizards is presented each time you add a new Tivoli Storage Manager server instance. Configure the Tivoli Storage Manager server to use the storage devices that you have available for client data, and set up the policies that you want client nodes to use.

The initial configuration process configures a single server. If you have purchased the enterprise edition administration feature and plan to configure a network of servers, you must perform additional tasks.

If you intend to configure Tivoli Storage Manager for use in a Microsoft Cluster Server (MSCS) environment, you must complete certain tasks before you begin the initial configuration of the Tivoli Storage Manager server.

> **Note:** You can click **Cancel** to exit any wizard panel. A dialog is displayed, asking whether you want to mark the current wizard task as complete.
>
> You can click **Yes** to continue to the next wizard, or click **No** to exit the initial configuration process. However, canceling during initial configuration can produce unexpected results.
>
> The preferred method is to complete the entire wizard sequence, and then restart an individual wizard to make any configuration changes.

You can choose from the following two wizard-based configuration paths: standard and minimal:

- Standard configuration

  Choose the standard configuration option to initialize and configure a server. A series of wizards is presented in sequence to guide you through the initial configuration process. We recommend this configuration path for setting up a functional production environment.

  During the standard configuration process, wizards help you perform the following commonly required tasks:

  - Analyze drive performance to determine the best location for Tivoli Storage Manager server components
  - Initialize the Tivoli Storage Manager server
  - Apply Tivoli Storage Manager licenses
  - Configure Tivoli Storage Manager to access storage devices

– Prepare media for use with Tivoli Storage Manager
– Register Tivoli Storage Manager client nodes
– Define schedules to automate Tivoli Storage Manager client tasks

Additional configuration wizards can help you perform these optional tasks:

– Configure Tivoli Storage Manager for use in a MSCS environment
– Configure Tivoli Storage Manager for use in a Microsoft Windows registry Active Directory environment
– Create a remote Tivoli Storage Manager for a Windows client configuration package

The standard configuration process does not include all Tivoli Storage Manager features, but it does produce a functional Tivoli Storage Manager system that can you can customize and tune further. The default settings suggested by the wizards are appropriate for use in many cases.

- Minimal configuration

  Choose this option to quickly initialize a Tivoli Storage Manager server instance and perform a test backup of data located on the Tivoli Storage Manager server machine. This configuration enables you to evaluate basic functions quickly.

  During the minimal configuration process, a wizard helps you initialize a Tivoli Storage Manager server instance. Open client registration is enabled, so Tivoli Storage Manager client nodes can automatically register themselves with the server. The following objects are also created on the server machine:

  – A client options file

    If a Tivoli Storage Manager client is not installed locally, the required directory structure is created. If a client options file already exists, it is backed up before the new file is created. TCP/IP communication is enabled for the client and server.

  – A file device

    A file device is drive space, which is designated for use as a virtual storage device. Standard files are used to represent individual media volumes. Data is written to file volumes sequentially, as though they were tape volumes. When a new file volume is required, a 25 MB file is automatically created. When file volumes are emptied, they are automatically deleted. Because the minimal configuration option does not provide for storage device configuration, default backup and archive storage pools are configured to send their data to the file device.

Although the wizards simplify the configuration process by hiding some of the details, configuring Tivoli Storage Manager for a Windows server requires a certain amount of Tivoli Storage Manager knowledge in order to create and maintain a typically complex storage management environment.

If you are planning to set up Tivoli Storage Manager for immediate use in a production environment, see Appendix A, "Performing a minimal configuration," of the *IBM Tivoli Storage Manager for Windows Installation Guide*, GC32-1602, for more information about configuration and management wizards.

You can use the command-line interface to perform all Tivoli Storage Manager configuration and management tasks, but the wizards are the preferred method for initial configuration. You can return to individual wizards after the initial configuration to update settings and perform management tasks. Again, you can refer to *IBM Tivoli Storage Manager for Windows Installation Guide*, GC32-1602, for more information about configuration and management wizards.

This section contains an overview of the wizard-based initial configuration process and instructions for performing the initial configuration. After you have installed Tivoli Storage Manager, complete the following steps:

1. Double-click the **Tivoli Storage Manager Management Console** icon on the desktop to open the Tivoli Storage Manager console window, shown in Figure 4-6.

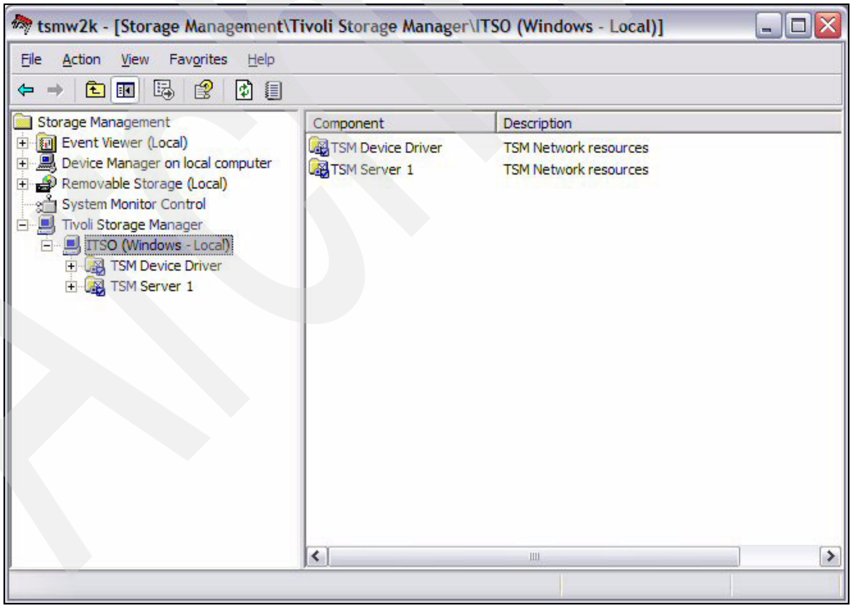

Figure 4-6 Tivoli Storage Manager management console

2. Expand the IBM Tivoli Storage Manager tree in the left pane until the local machine name is displayed.
3. Right-click the local machine name, and select **Add a New Tivoli Storage Manager Server**.

> **Note:** You can also use the Health Monitor to add a new Tivoli Storage Manager server.

4. The Initial Configuration Task List is displayed. Select **Standard configuration** or **Minimal configuration**, and click **Start**.

> **Note:** If a Tivoli Storage Manager server instance already exists on the local machine, you are prompted to confirm that you want to create and configure a new server instance. Be careful to create only the server instances you require. In most cases, only one server instance is necessary.

Various wizards are displayed sequentially to enable you to configure the Tivoli Storage Manager server instance. We discuss the following wizards in this section:

– "Running the Initial Configuration Environment Wizard" on page 131
– "Running the Performance Configuration Wizard" on page 133
– "Running the Server Initialization Wizard" on page 135
– "Running the Device Configuration Wizard" on page 137
– "Running the Client Node Configuration Wizard" on page 139
– "Running the Media Labeling Wizard" on page 144
– "Completing initial configuration" on page 147

Running the Initial Configuration Environment Wizard

The Initial Configuration Environment Wizard is the first wizard in the standard configuration sequence. The information you provide in this wizard is used to customize upcoming wizards to reflect your preferences and storage environment.

This wizard consists of a welcome page and a series of input pages that help you perform the following tasks:

- Choose whether configuration tips are automatically displayed during the initial configuration process. This additional information can be helpful for new Tivoli Storage Manager users.
- Choose to configure Tivoli Storage Manager in a standalone or network environment. Table 4-1 describes these environments.

Table 4-1 Environment description

Tivoli Storage Manager environment	**Description**
Stand-alone	A Tivoli Storage Manager backup/archive client and Tivoli Storage Manager server are installed on the same machine to provide storage management for this machine only. No network-connected Tivoli Storage Manager clients are installed. Client/server communication is configured automatically.
Network	A Tivoli Storage Manager server is installed. The backup/archive client is optionally installed on the same machine. You are licensed to install network-connected Tivoli Storage Manager clients on remote machines. You must configure communications between the remote clients and the server.

Running the Performance Configuration Wizard

The Performance Configuration Wizard consists of a welcome page and a series of input pages that prompt you for input. Complete the following steps to run the wizard:

1. Estimate the number of clients the Tivoli Storage Manager server supports and the typical size of files to be stored (see Figure 4-7). Choose the appropriate number of clients and estimated file sizes in your environment. Click **Next**.

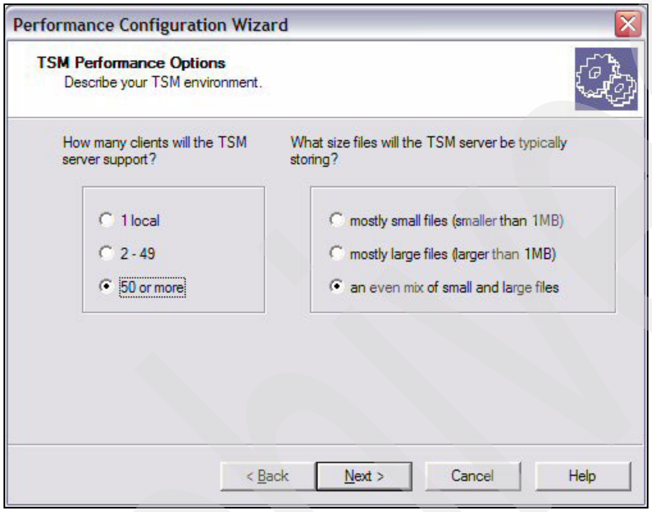

Figure 4-7 Describing the Tivoli Storage Manager environment

2. Tivoli Storage Manager analyzes local drives to determine the best location for initial Tivoli Storage Manager server volumes (see Figure 4-8).

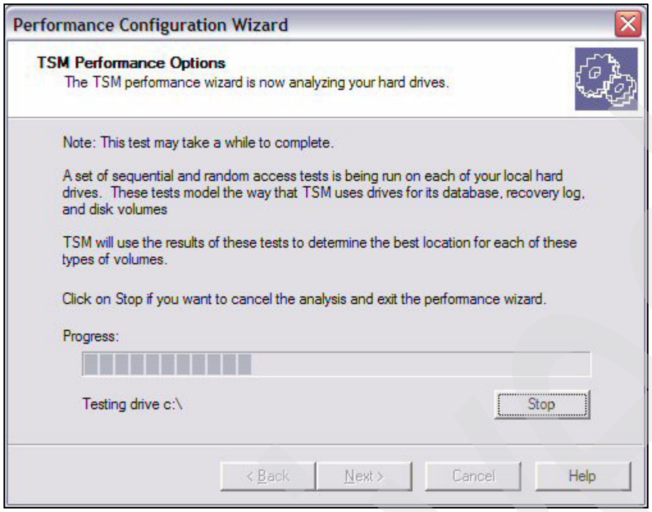

Figure 4-8 Searching for the best drive

The information you provide in this wizard and the results of an automated analysis of the local drives are used to determine the best location for three important Tivoli Storage Manager volumes (see Table 4-2). The term *volume* is used here to refer to space allocated in server disk storage.

Table 4-2 Overview of volumes

Volume	Description
Database	Stores information required for server operations and information about client data that has been backed up or archived.
Recovery log	Stores information about pending database updates until they are committed.
Disk storage pool	Stores client data in server disk space. This data can be kept on the server or migrated to another storage resource.

Preferred locations for these Tivoli Storage Manager volumes are displayed in the Server Initialization Wizard; see “Running the Server Initialization Wizard” on page 135 for modifying the volume locations and default sizes.

Running the Server Initialization Wizard

In this section we describe the steps to initialize Tivoli Storage Manager servers. The Server Initialization Wizard is the only wizard that is displayed during the minimal configuration process. It also is displayed as part of the standard configuration wizard sequence. This wizard consists of a welcome page and a series of input pages that help you perform the following tasks:

1. Choose a directory to store files unique to the Tivoli Storage Manager server instance that you are currently configuring (see Figure 4-9). Click **Next**.

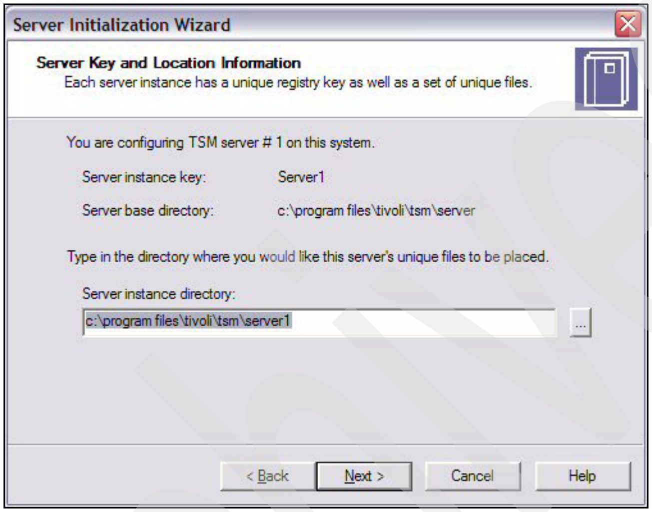

Figure 4-9 Starting server initialization

2. Choose directories for initial database, recovery log, and disk storage pool volumes (see Figure 4-10 on page 136). The default locations are preferred, based on the results of the Performance Configuration Wizard analysis. Choose whether to dynamically extend the size of the database and recovery log by adding volumes as necessary. Click **Next**.

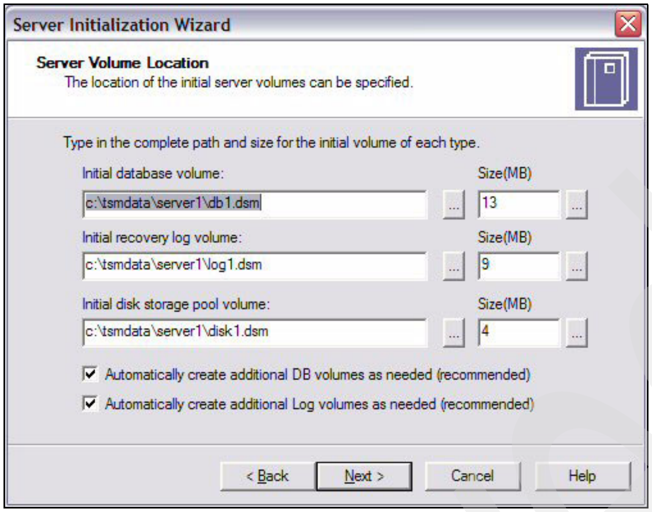

Figure 4-10 Database configuration options

3. Choose a logon account for the Tivoli Storage Manager server service, and choose whether the service is started manually or automatically.
4. Choose a name and password for the Tivoli Storage Manager server. Some Tivoli Storage Manager features require a server password.
5. If a Microsoft cluster server is detected during the standard configuration process, you are prompted to configure Tivoli Storage Manager for use in a clustered environment. Select **Yes** to start the Cluster Configuration Wizard. Before you set up a cluster for use with Tivoli Storage Manager, you must do some planning and ensure that your hardware is supported.

> **Note:** The minimal configuration process does not support cluster configuration.

When you complete the Server Initialization Wizard, Tivoli Storage Manager:

- Initializes the server database and recovery log.
- Creates the database, recovery log, and disk storage pool initial volumes.

  Server Initialization Wizard default settings specify the following initial volume sizes (by default, additional volumes are dynamically added to extend the database and recovery log as required):

  – A 13 MB database volume (db1.dsm).

  – The database size is largely determined by the amount of client data managed by the server. As you add clients, the database size may increase.

- A 9 MB recovery log volume (log1.dsm). The recovery log size depends on the database size, and increases accordingly.
- A 4 MB storage pool volume (data1.dsm).

- Creates two default schedules: DAILY_INCR and WEEKLY_INCR. You can use the Scheduling Wizard to work with these schedules or create others.
- Registers a local administrative client with the server. This client is used to provide access to the administrative Web interface and server command-line interface. The client is named admin, and its default password is admin. To ensure system security, we recommend that you change this password.

Initialization results are recorded in the initserv.log file in the server directory. If you have problems starting the server after initialization, check this log file for error statements. If you contact technical support for help, you might be asked to provide this file.

Running the Device Configuration Wizard

The Device Configuration Wizard (shown in Figure 4-11) automatically detects storage devices attached to the Tivoli Storage Manager server. Use this wizard to select the devices that you want to use with Tivoli Storage Manager and configure device sharing if required.

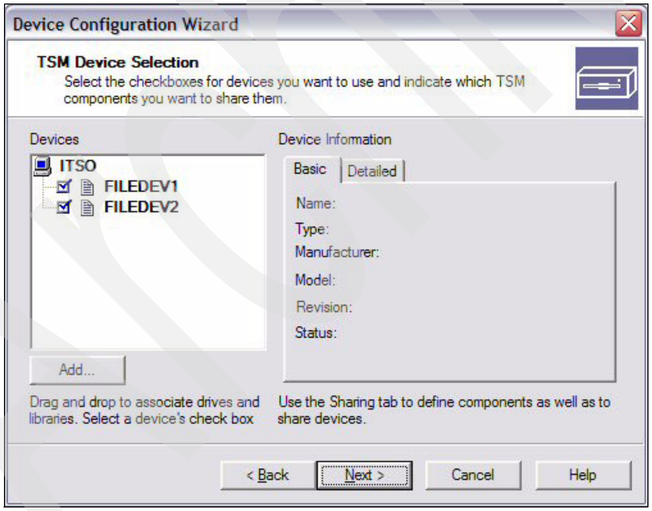

Figure 4-11 Device Configuration Wizard

The wizard pane on the left displays a tree view of devices connected to the Tivoli Storage Manager server machine. Tivoli Storage Manager device names are used to identify devices. Libraries and drives can be detected only if your hardware supports this function.

The right pane in Figure 4-11 on page 137 displays basic and detailed information about the device selected in the tree view. If the device is a type that can be shared, the Sharing tab displays any Tivoli Storage Manager components that share the device.

To define a device, select its check box. Any device with an open check box can be defined to the Tivoli Storage Manager server. A library check box that is partially filled indicates that some of the drives associated with that library have not been selected for use with Tivoli Storage Manager.

Note: A solid green check box indicates that the device was previously defined on Tivoli Storage Manager. Previously defined devices cannot be manipulated or removed using the wizard. You can use the administrative Web interface or server command line to perform this task.

The Device Configuration Wizard consists of a welcome page and input pages that help you perform the following tasks:

- Select the storage devices that you want to use with Tivoli Storage Manager, and define them
- Manually associate drives with libraries, if required
- Specify the Small Computer System Interface (SCSI) element number order for manually associated drives
- Configure device sharing, if required
- Manually add virtual or undetected devices
- Manually associate drives

You must manually associate any drive listed as Unknown with a library. For example, drives attached to a Fibre Channel switch or a storage area network (SAN) cannot be automatically associated. Tivoli Storage Manager can determine that the library contains a certain number of drives but cannot acquire their element numbers or addresses. The correct names for these drives appear at the bottom of the tree as standalone drives. Drag and drop the unknown drive on the correct library. To use a library with Tivoli Storage Manager, you must replace the names of any of drives displayed as Unknown with a valid drive name.

Note: If you manually associate more than one drive with the same library, you must order the drives according to element number. If you do not arrange the drives correctly, Tivoli Storage Manager does not work as expected. To determine the element number for a drive, select the drive and click the **Detailed** tab in the right wizard pane. Use the element number lookup tool to determine the correct position of the drive. If your drive is not listed, refer to the manufacturer's documentation.

- You can use the Device Configuration Wizards to set up device sharing. To do so, click the **Sharing** tab and click **Components**. Follow the directions in the Device Sharing dialog.
- The Device Configuration Wizard also enables you to add virtual or undetected devices. Click **New** to add file-type devices and drives or libraries accessed through an NDMP file server.

The libraries and drives you define to Tivoli Storage Manager are available for storing data.

Running the Client Node Configuration Wizard

With the Client Node Configuration Wizard (shown in Figure 4-12), you can add and register the client nodes that back up data to the server instance that you are currently configuring. The wizard also enables you to specify how the backup data for these clients is stored by associating client nodes with storage pools.

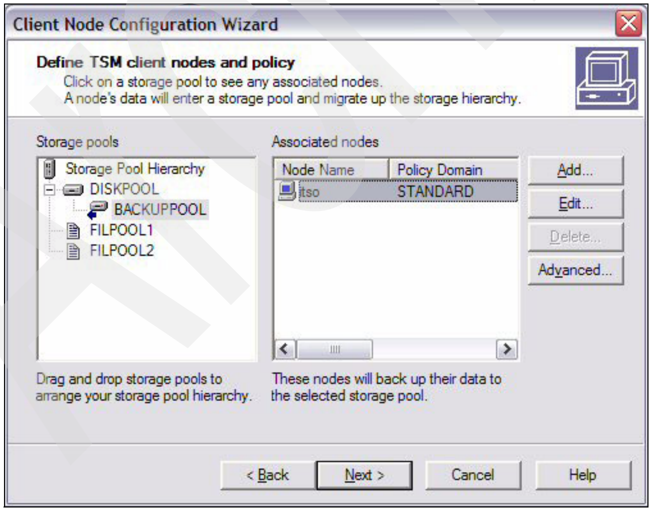

Figure 4-12 Client Node Configuration Wizard

The left pane of the wizard shown in Figure 4-12 on page 139 displays two default Tivoli Storage Manager storage pools (DISKPOOL and BACKUPPOOL). If you used the Device Configuration Wizard to define any storage devices to Tivoli Storage Manager, storage pools associated with these devices are automatically generated and are also displayed in the left pane of the Client Nodes Configuration Wizard.

Tivoli Storage Manager uses a logical construct called a *storage pool* to represent storage resources. Different storage pools are used to route client data to different kinds of storage resources. Storage pools can be arranged in a hierarchy, with one pointing to another, to allow for migration of data from one type of storage to another.

Tivoli Storage Manager provides a default storage pool named DISKPOOL, which represents random access storage space on the hard drive of the Tivoli Storage Manager server machine. During server initialization, Tivoli Storage Manager created one volume (representing a discrete amount of allocated space) in this storage pool. By default, this volume was configured to grow dynamically. You can add more volumes to expand this storage pool as required.

Tivoli Storage Manager also provides three other default storage pools, which are set up to point to DISKPOOL. These three storage pools correspond to the three ways Tivoli Storage Manager manages client data (backup, archive, and space management). The Client Node Configuration Wizard enables you to work with the backup storage pool BACKUPPOOL.

By default, data for any client nodes you associate with BACKUPPOOL are transferred immediately to DISKPOOL. You can store the data in DISKPOOL indefinitely, or just use DISKPOOL as a temporary cache and then migrate the data to any other storage devices represented in the storage pool hierarchy.

The right pane displays client nodes associated with the storage pool selected in the left pane.

To register new client nodes, you must provide client node names and passwords. You can also change storage policy settings by adding or modifying policy domains. Tivoli Storage Manager storage policy determines how many copies of backed up files are maintained, and how long individual copies of files are retained in storage.

> **Note:** Consider using this wizard to register any remote client nodes at this stage, even if you have not yet installed Tivoli Storage Manager Client code on these machines. After you complete the initial server configuration, you can install the client code and configure the client nodes to start transferring data to this server.

The Client Node Configuration Wizard consists of a welcome page and several input pages that help you perform the following tasks:

- Register client nodes with the Tivoli Storage Manager server. You can add nodes individually, or detect and register multiple clients at once. To register client nodes individually, perform the following tasks:

  a. Click **Add**. The Properties dialog is displayed, which is open at the Node information tab, as shown in Figure 4-13.

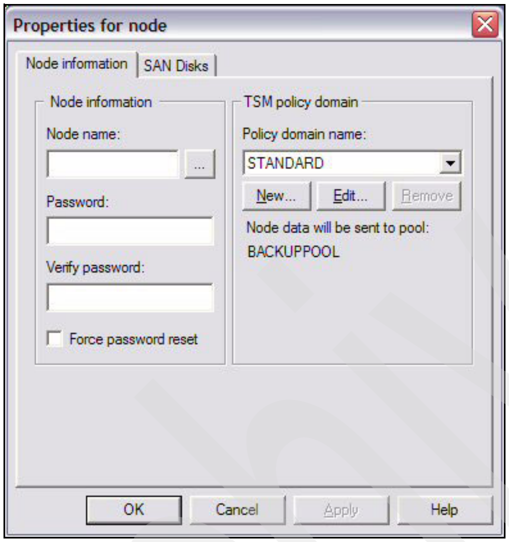

Figure 4-13 Node properties

  b. Enter the node name and password information.

  c. Consider your storage policy requirements. By default, the new client node is associated with the STANDARD storage policy domain. BACKUPPOOL is the default backup storage pool for this domain. You can associate the new client node with a different storage pool by clicking **New** to create a new policy domain or by clicking **Edit** to modify the existing policy domain.

     Managing multiple policy domains can significantly increase your administrative overhead, so you must create only the domains you require.

  d. To detect and register multiple client nodes at once, return to the main wizard panel and click **Advanced**. Follow the instructions in the Properties dialog. You can add clients from a text file or choose from computers detected in your Windows domain. The Tivoli Storage Manager console directory contains a file named sample_import_nodes.txt, which defines the format required to import client nodes.

e. To modify Tivoli Storage Manager client node information, select a client node name from the right wizard pane and click **Edit**. To delete the client node you just added, select the client node name and click **Delete**.

Note: You cannot use the wizard to delete a client that was previously defined to the server. You can use the administrative Web interface or server command line to perform this task.

- Define client/disk associations for use withTivoli Storage Manager server-free data movement.

  You can configure Tivoli Storage Manager for server-free data movement, which allows for full-volume backup and restore of client data stored on a dedicated SAN-attached disk. If you are planning to use this feature, you must define the client/disk association to the Tivoli Storage Manager server:

  a. From the main wizard panel, click **Add**. The Properties dialog is displayed.

  b. Click the **SAN Disks** tab (see Figure 4-14). The Server Free Data Movement Disk Information page is displayed.

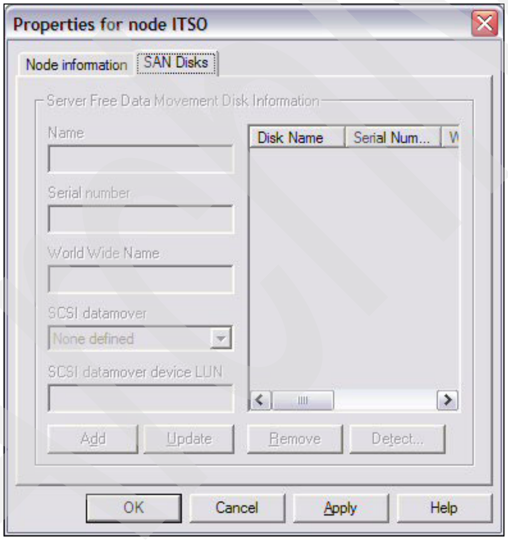

Figure 4-14 Configuring disk information

c. To manually add SAN disk information:

i. Enter the name that uniquely identifies this SAN disk. Use the format *harddiskX*, where *X* is the disk number defined to the client machine. You can use the Microsoft Management Console (MMC) Disk Management snap-in to obtain this disk number from the client machine (see Tivoli Storage Manager online help):

`http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp`

ii. Enter the serial number that identifies this disk on the SAN. You can obtain the serial number from the disk.

iii. Optionally, enter a global name for the disk.

iv. Click **Add**.

The disk is added to the right pane (see Figure 4-14 on page 142) and is associated with this client node when the wizard completes. The following tasks are also related to SAN disks:

- To quickly identify and add SAN disk information, click **Detect** and follow the instructions in the Detect SAN Devices dialog.
- To modify SAN disk information, select the disk name in the right pane, update any fields, and click **Update**.
- To remove a SAN disk from the list, select the disk name in the right pane, and click **Remove**.

Note: You cannot use the wizard to remove a disk that was previously associated with this client. You can use the administrative Web interface or server command line to perform this task.

- Arrange the storage pool hierarchy to meet your storage requirements.

By default, new client nodes send backup data to BACKUPPOOL, which immediately migrates the data to DISKPOOL. You can point BACKUPPOOL to any other displayed storage pool to route data there instead. A storage pool can migrate data to one other storage pool. Multiple storage pools can be set up to migrate data to the same storage pool.

To see which clients are associated with a storage pool, select a storage pool in the left wizard pane. Any client nodes associated with this pool are displayed in the right pane.

Note: In a standalone server configuration, it is generally more efficient to back up data directly to tape. However, in a network configuration, consider arranging your storage pools so that client data is backed up to disk and later migrated to tape.

- Associate registered nodes with storage pools by adding the clients to a new or existing policy domain. For example, these steps describe how to back up data directly to tape or to disk. To back up client data directly to tape, complete the following steps:

  a. Associate clients with BACKUPPOOL.
  b. Drop BACKUPPOOL on a tape storage pool (for example, 8MMPOOL1).

  To back up client data to disk, for migration to tape, follow these steps:

  a. Associate clients with BACKUPPOOL.
  b. Drop BACKUPPOOL on DISKPOOL (the default setting).
  c. Drop DISKPOOL on a tape storage pool.

Client nodes that you have registered can be configured to back up data to this Tivoli Storage Manager server instance. The backup data is managed according to the way that you set up the client's associated storage pool hierarchy.

Running the Media Labeling Wizard

Storage media must be labeled and checked into Tivoli Storage Manager before it can be used. Media labels are written at the start of each volume to uniquely identify this volume to Tivoli Storage Manager.

The Media Labeling Wizard is displayed only if you have defined attached storage devices to Tivoli Storage Manager. Slightly different versions of the wizard are displayed for automated and manual storage devices. This section describes the media labeling and check-in process for automated library devices.

The Media Labeling Wizard consists of a welcome page and a series of input pages that help you perform these tasks:

1. Select the devices that contain the media that you want to label.
2. Select and label specific media.
3. Check in labeled media to Tivoli Storage Manager.

In the window shown in Figure 4-15, you can choose devices for labeling purposes. The left pane displays devices and drives recognized by Tivoli Storage Manager. The right pane displays information about any device or drive selected in the left pane. To select a device and any associated drives, check the box next to the device or drive name.

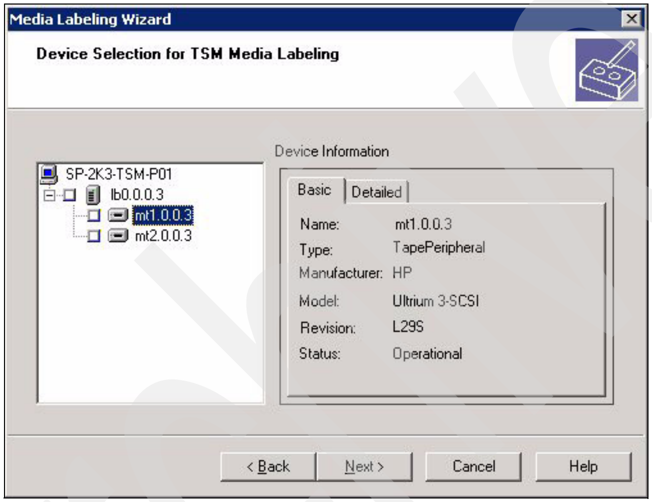

Figure 4-15 Tivoli Storage Manager device selection

To select and label media, complete the following steps:

1. Check the box next to the media that you want to label (see Figure 4-16).
2. Check **Overwrite existing label** if necessary, and select from the other available labeling options.
3. Click **Label Now**.

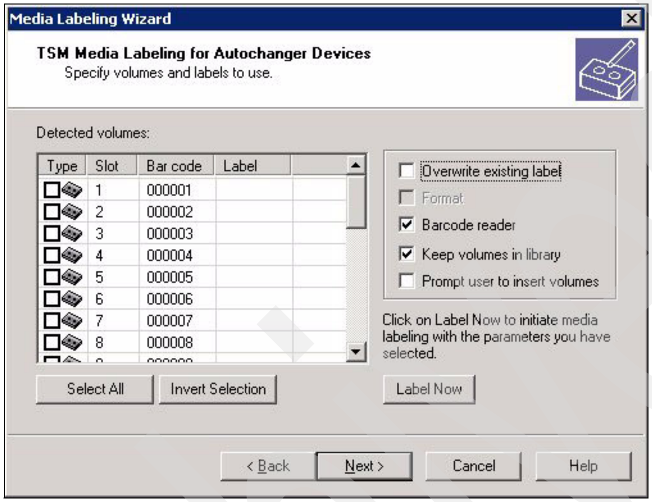

Figure 4-16 Media Labeling Wizard

4. The Tivoli Storage Manager Media Labeling dialog is displayed. Enter a label for the media. The Media Labeling Wizard supports labels that are up to six characters long.
5. Click **OK**. The Tivoli Storage Manager Media Labeling Monitor dialog is displayed. The status is displayed and updated throughout the labeling process.

   When the labeling process is complete, the OK button becomes active. The amount of time this takes can depend on the storage hardware and type of media you are using.
6. Click **OK**. The new label is displayed in the left pane.
7. After you have finished labeling media, click **Next**. The Media Check-in dialog is displayed.

   The Media Check-in dialog is displayed only if you labeled media in the previous dialog.

8. To check in labeled media to Tivoli Storage Manager, click **Check-in now**.

   Media volumes from all of the storage devices you selected in the first media labeling dialog are eligible for check-in. All labeled media not previously checked in to this server are checked in automatically at this time. A dialog is displayed, describing the check-in process. Check-in runs as a background process, and media is available for use until the process completes.

Depending on your storage hardware and the amount of media being checked in, this process can take some time. To monitor the check-in process, complete the initial configuration and follow these steps:

1. From the Tivoli Storage Manager Console, expand the tree for the Tivoli Storage Manager server you are configuring.
2. Expand **Reports** and click **Monitor**.
3. Click **Start** to monitor server processes in real time.

When the check-in process has completed, media is available for use by Tivoli Storage Manager. By default, media volumes are checked in with scratch status.

Completing initial configuration

After the initial configuration completes, you are prompted to verify your configuration. If you have installed a local backup/archive client, click **Yes** to start the client immediately. Click **No** if you have not installed the client code locally, or if you plan to verify your configuration by backing up remotely installed clients.

> **Note:** Click the **Tivoli Storage Manager Backup Client** icon on your desktop to start the local backup/archive client at any time. You can use the Tivoli Storage Manager console to perform a variety of administrative tasks, including issuing commands and monitoring server processes. You can also access the individual wizards you used during the initial configuration process from this interface. Additional wizards are also available.

The Tivoli Storage Manager configuration wizards simplify the setup process by hiding some of the detail. For the ongoing management of your Tivoli Storage Manager system, it can be helpful to understand what has been created for you.

Table 4-3 lists the default Tivoli Storage Manager data management policy objects.

Table 4-3 Default data management policy objects

Tivoli Storage Manager objects	Name	Details
Policy Set	STANDARD	This set is active. It contains one management class.
Management Class	STANDARD	Contains a backup copy group and an archive copy group.
Copy Group (Backup)	STANDARD	Stores one active and one inactive version of existing files. The inactive version is kept for 30 days. Stores one inactive version of deleted files for 60 days. Points to BACKUPOOL.
Copy Group (Archive)	STANDARD	Stores one active and one inactive version of existing files. The inactive version is kept for 30 days. Stores one inactive version of deleted files for 60 days. Points to ARCHIVEPOOL.

After installing the Version 5.5.0 server, you must register new licenses. The REGISTER LICENSE command enables licenses for certain components:

- You can register licenses for server components. This includes Tivoli Storage Manager (base), Tivoli Storage Manager Extended Edition, and Tivoli Data Retention Protection.
- You cannot register licenses for components that are licensed on the basis of processors (for example, Tivoli Storage Manager for Mail, Tivoli Storage Manager for Databases, Tivoli Storage Manager for ERP, Tivoli Storage Manager for Hardware, and Tivoli Storage Manager for Space Management).

Your license agreement determines what you are licensed to use, even if you cannot use the REGISTER LICENSE command to register all components. You are expected to comply with the license agreement and use only what you have purchased. Use of the REGISTER LICENSE command implies that you agree to and accept the license terms specified in your license agreement. The command syntax is:

```
REGister LICense FILE=<license_file>
```

You can use one of the following file parameters:

tsmbasic.lic — To license base IBM Tivoli Storage Manager.

tsmee.lic — To license base IBM Tivoli Storage Manager Extended Edition. This includes the Disaster Recovery Manager (DRM), large libraries, and Network Data Management Protocol (NDMP).

dataret.lic	To license IBM Tivoli Storage Manager for Data Retention. This parameter is required to enable Tivoli Data Retention Protection and Expiration and Deletion Suspension (Deletion Hold).

4.2.5 Installing the IBM Tivoli Storage Manager backup/archive client

It is always a best practice to install the Tivoli Storage Manager backup/archive client on the Tivoli Storage Manager server machine. This helps you to validate the Tivoli Storage Manager server you configured. This section describes the steps you must complete to install the Tivoli Storage Manager backup/archive client. For more information, see *IBM Tivoli Storage Manager for Windows: Backup-Archive Clients Installation and User Guide Version 5.5*, SC32-0146-01. This document is available from the following Web site:

http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/topic/com.ibm.itsmfdt.doc/ans60000.htm

The Tivoli Storage Manager backup/archive client helps you protect information on your workstation. While using Tivoli Storage Manager, you can maintain backup versions of your workstation files that you can restore if the original files are damaged or lost. You can also archive workstation or server files that you do not currently require, preserve them in their current state, and retrieve them when necessary.

You can access Tivoli Storage Manager backup and archive features:

- Locally through the client command line or GUI interface
- Remotely or locally through the Web client interface

Note: If a client upgrade is performed, you must reboot the system if any instance of the client is running during the install. Stop all instances of the Tivoli Storage Manager client (services, interactive clients, and so on) before performing the install. If the Logical Volume Snapshot Agent (LVSA) component is selected for install, then you must reboot in order to install or update the tsmlvsa.sys filter driver.

Complete the following steps to install the software on your Microsoft Windows XP Professional SP2 or higher, Windows Server 2003 (all editions, 32 and 64 bit), Windows Server 2003 R2 (all editions, 32 and 64 bit), and Windows Vista® (all editions, 32-bit and 64 bit) workstations:

1. Insert the CD that contains the Tivoli Storage Manager Windows client into your CD drive. If you have autorun enabled, the installation dialog starts when the CD loads. If the installation dialog does not start automatically, you can start it manually. Select **Run** from the Start menu, and at the prompt, type `x:\setup` (where `x` is your CD drive). Click **OK**.

> **Note:** If you do not have the CD and downloaded the software from the Passport Advantage Web site, simply execute the .exe file. The file is self-extractable, and you are directed to the installation window shown in Figure 4-17.

2. First, you must choose your language for your menu selections, and then you access the installation main menu.
3. The Tivoli Storage Manager Client InstallShield Wizard welcome window is displayed, as shown in Figure 4-17. Choose **Next** to continue.

Figure 4-17 IBM Tivoli Storage Manager Client installation main menu

4. Choose the destination folder for Tivoli Storage Manager Client code and click **Next** to continue. See Figure 4-18.

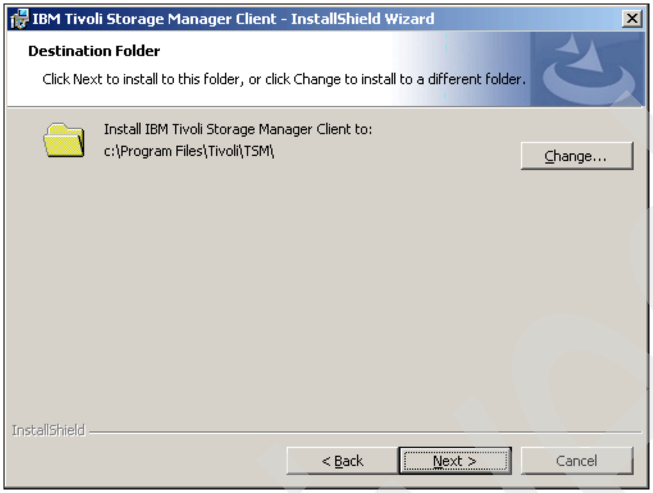

Figure 4-18 Choosing destination folder menu

You access the Setup Type option menu.

5. Choosing **Typical** installs the minimum necessary to provide normal backup and archive functions. If you are upgrading from a previous version, this selection updates the currently installed components. Otherwise, it installs the backup/archive client, the API Runtime files, and the Web client.

6. At this point, choose **Custom** and click **Next**.

   Choosing **Custom** takes you to the Custom Setup window. From this window, you can click any program feature icon to modify this feature if it is not mandatory for installation.

   You can select from the following program features:

   – Backup/archive client:

     - Backup/archive client GUI files (mandatory; cannot be deselected)
     - Backup/archive client Web files (mandatory; cannot be deselected)
     - Client application programming interface (API) runtime files (mandatory; cannot be deselected)
     - Client API Software Developer's Kit (SDK) files (optional)
     - Administrative client command-line files (optional; not enabled by default)
     - Logical Volume Snapshot Agent (LVSA) (optional; not enabled by default)

– Online information:

- READMEs and *Installation and User's Guide:*

  http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp

- *Installation and User's Guide - PDF* (optional)
- Language support

> **Attention:** The Tivoli Storage Manager client now makes use of language packs for non-English language support. Each supported language has its own installation package that must be installed in order to use Tivoli Storage Manager in a supported, non-English language. The Tivoli Storage Manager client is a prerequisite for installing a Tivoli Storage Manager client language pack.

7. Click **Next** to start the installation (see Figure 4-19).

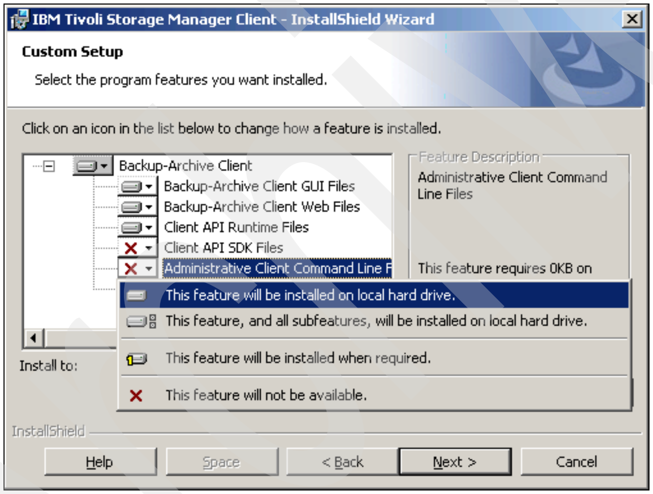

Figure 4-19 Client custom setup options

8. The ready to install window is displayed. Click **Install** to proceed.
9. When you receive the InstallShield Wizard Completed message, click **Finish** to complete the installation.

4.2.6 Configuring the IBM Tivoli Storage Manager client

Configure the communications options in the client options file to connect to the server. Each client requires a client options file, which contains options that identify the server, communication method, backup and archive options, space management options, and scheduling options.

You can edit or create client options files in several ways, depending on the client platform and configuration of your system:

- Any client

  Edit the dsm.opt client options file with a text editor at a client workstation. This method is the most direct, but it might not be best if you have many clients.

- Microsoft Windows clients

  Generate the dsm.opt client options file from the server with the Network Client Options File Wizard. This method is easy and direct, and the wizard detects the network address of the Tivoli Storage Manager server. To run the wizard, complete these steps:

  a. From the Tivoli Storage Manager console, expand the tree for the Tivoli Storage Manager server on which you want to create the file and click **Wizards**. The Wizards list is displayed in the right pane.

  b. Double-click **Client Options File** from the Wizards list to start the wizard.

  c. Follow the instructions in the wizard.

- Networked Windows clients with a shared directory on a file server

  Use the Remote Client Configuration Wizard to create a package that enables remote users to create client options files. The administrator uses the wizard to generate a client configuration file and stores the file in a shared directory. Clients access the shared directory and run the configuration file to create the client options file. This method is suitable for sites with many clients.

Each Tivoli Storage Manager client instance requires a client options file (dsm.opt). In a UNIX environment, you require both the dsm.sys and dsm.opt files. The minimum basic communication option for the dsm.opt file is shown in Example 4-1.

Example 4-1 Minimum dsm.opt file

```
commmethod       TCPIP
tcpport          1500
tcpserveraddress localhost
```

For more information about configuring the client, refer to the *Backup-Archive Clients Installation and User's Guide* that corresponds to your operating system.

4.2.7 Backing up a client

This section describes the process of completing simple backup of client files in a Microsoft Windows environment.

Note: We recommend that you back up a small file or directory.

Perform the following steps to back up a remote or local client:

1. Start the client, enter a node name and password, and click **Login**.

   Once the client is authenticated with the server, the backup/archive client window opens, as shown in Figure 4-20.

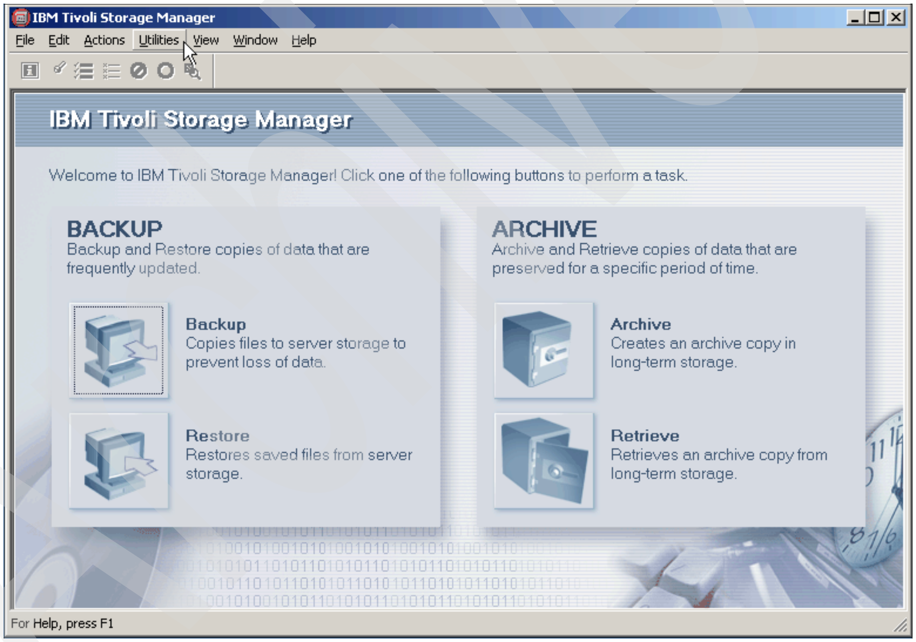

Figure 4-20 Backup/archive client GUI

2. Click **Backup** from the client window. The Backup window opens.
3. Expand the directory tree, shown in Figure 4-21

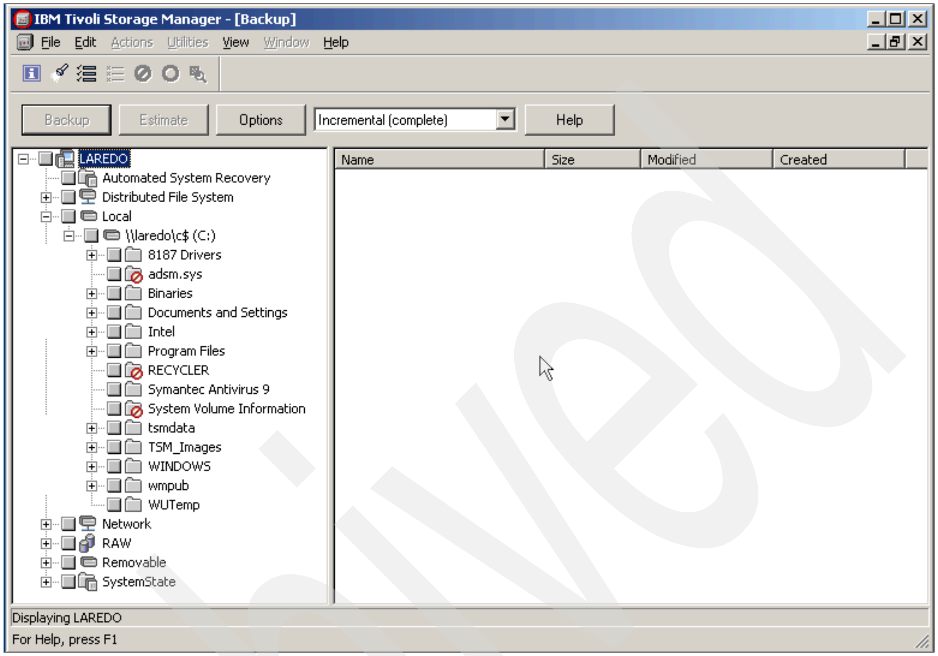

Figure 4-21 Expanded directory tree

4. Select the folder icons to display the files in the directory.
5. Check the selection boxes next to the files or directories you want to back up.

6. From the pull-down list shown in Figure 4-22, choose the backup type:
   – Incremental (complete)
   – Incremental (date only)
   – Always backup (for a selective backup)

Note: The first backup of a file is always a full backup, regardless of what you specify.

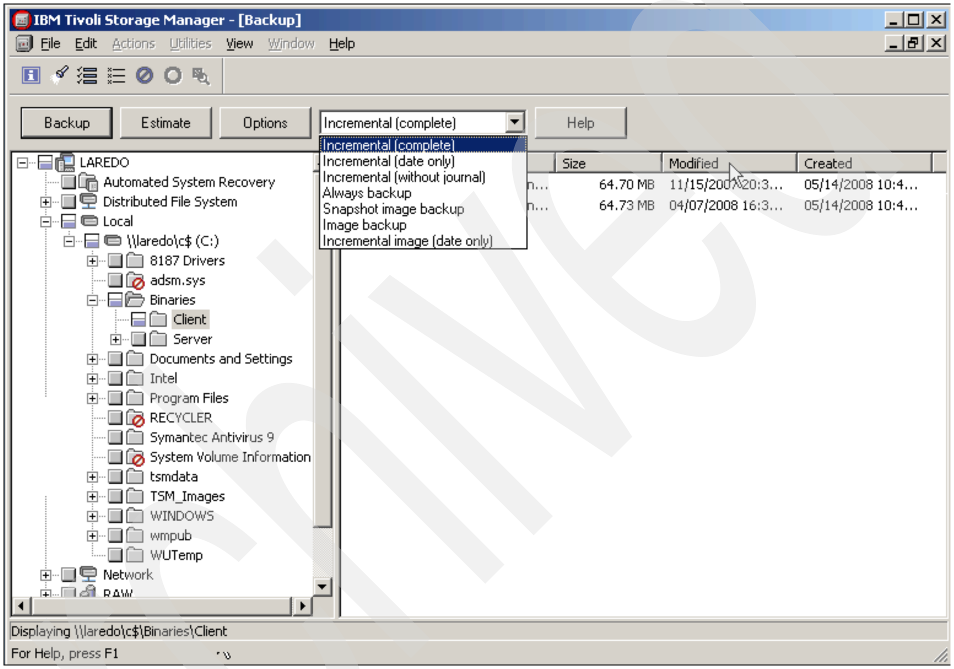

Figure 4-22 Backup selection

7. Click **Backup**. The Backup Report window, shown in Figure 4-23, displays the backup processing status.

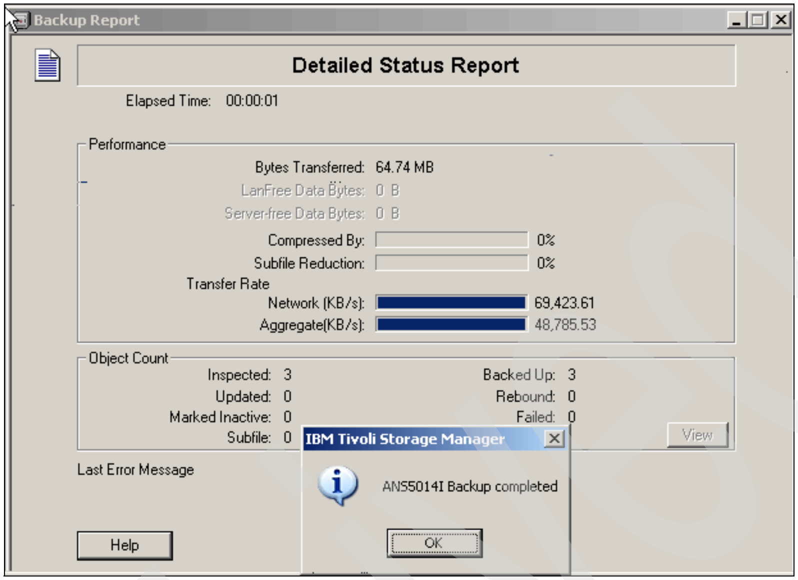

Figure 4-23 Incremental backup report

> **Note:** If you receive a message indicating that the server cannot be contacted, see the Common Questions and Answers section of the IBM Tivoli Storage Manager Online Information:
>
> http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp

Excluding files from the backup

You might not want to back up certain files that might be core files, local caches of network file systems, OS or application files that can be easily recovered by installing the program again, or any other files that you can easily rebuild. You can create an include-exclude list in the client options file to exclude certain files from both incremental and selective backup processing. Tivoli Storage Manager backs up any file that is not explicitly excluded from backup. You can also include specific files that are in a directory that you have excluded. For more information, see the *Backup/Archive Clients Installation and User's Guide* for your particular platform.

Restoring client files or directories

This section describes a simple restore of client files. For details and advanced procedures, see the appropriate *Backup/Archive Clients Installation and User's Guide* publication. To restore backup versions of files or directories, complete the following steps:

1. Click **Restore** from the client window. The Restore window opens.
2. Expand the directory tree.
3. Expand the File Level.
4. Click the selection boxes next to the files or directories that you want to restore.
5. Click **Restore**. The Restore Destination window opens.
6. Select the destination in the Restore Destination window.
7. Click **Restore**. The Restore Report window displays the restore processing status.

Tivoli Storage Manager can keep multiple versions of files, and you can choose which version to restore. Tivoli Storage Manager marks the most recent version as active and all other versions as inactive. When you back up a file, Tivoli Storage Manager marks the new backup version active and the previous active version as inactive. When the maximum number of inactive versions is reached, Tivoli Storage Manager deletes the oldest inactive version.

If you try to restore both an active and inactive version of a file at the same time, only the active version is restored. The following instructions specify how to restore an active and inactive backup version:

- To restore an active backup version, click **Display active files only** from the View pull-down list.
- To restore an inactive backup version, click **Display active/inactive files** from the View pull-down list.

Archiving and retrieving files

This section describes the process of a simple archive and retrieval of client files. For more information, see the appropriate *Backup/Archive Clients Installation and User's Guide* publication.

We recommend that you archive a small file or directory. You can select files to be archived by name or from a directory tree.

To archive files by name, perform the following steps:

1. Click **Archive** in the client main window. The Archive window opens.
2. Expand the directory tree until you find the drive or directory that you want.
3. Highlight the drive or directory that you want.
4. Search for file names:
   a. Click the **Find** icon on the tool bar.
   b. Enter the search criteria in the Find Files window, shown in Figure 4-24. You can use a mask to find files with similar names. Assign a unique description for each archive package.

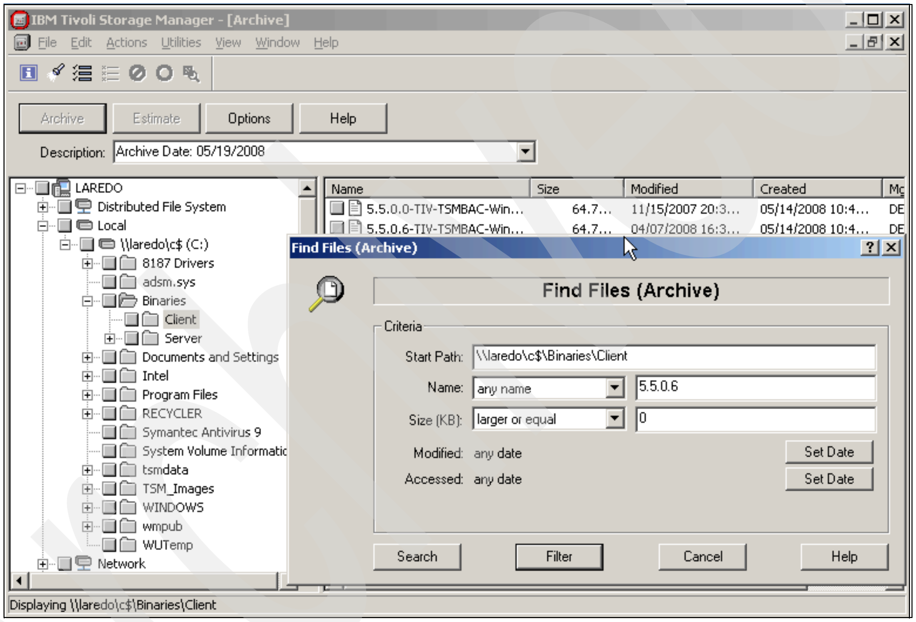

Figure 4-24 Archive using "find files"

   c. Click **Search**. The Matching Files window opens.

5. Click the selection boxes next to the files that you want to archive.
6. From the Window menu, select **Archive** to go back to the archive window.
7. In the Description box on the tool bar, enter a description, accept the default description, or select an existing description for your archive package.
8. Click **Archive** to archive the files. The Archive Status window displays the status progress of the archive.

Archiving files using a directory tree

You can archive specific files or entire directories from a directory tree. To archive your files from the directory tree, perform the following steps:

1. Click **Archive** in the client main window, shown in Figure 4-25. The Archive window opens.
2. Expand the directory tree until you find the directories or drive that you want.
3. Click the selection boxes next to the files or directories to archive.
4. In the Description box on the tool bar, enter a description, accept the default description, or select an existing description for your archive package.

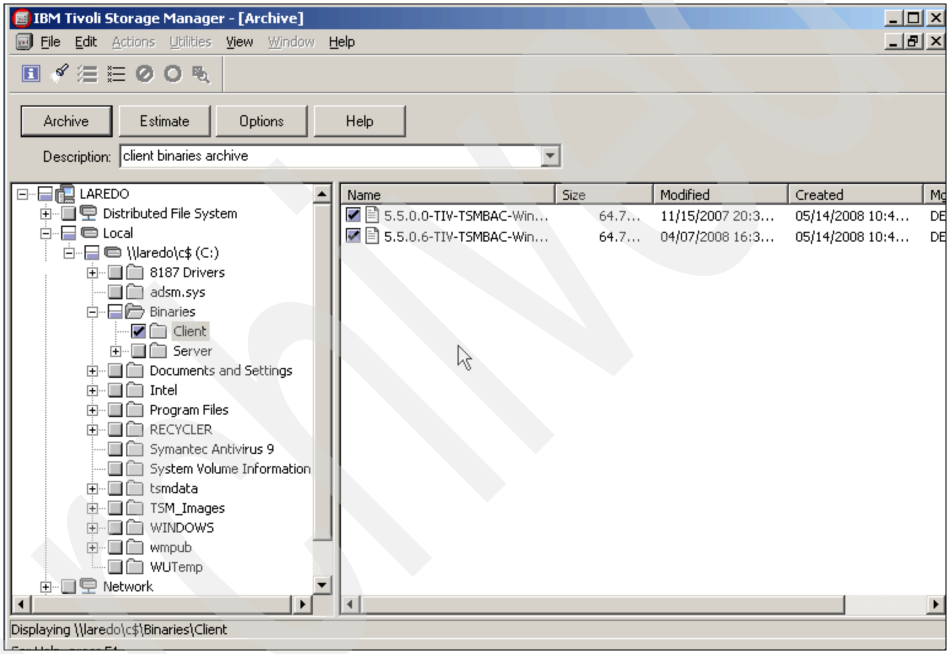

Figure 4-25 Archive using directory tree

5. Click **Archive**. The Archive Status window opens. The Archive Report window displays the status progress of the archive.

Retrieving archive copies

You retrieve files when you want to return archived copies of files or directories to your workstation. To retrieve archived copies, perform the following steps:

1. Click **Retrieve** on the client main window. The Retrieve window opens.
2. You can find the files or directories in either of the following ways:
   - From the directory tree:

     Expand the directory tree until you find the object that you want. The objects are grouped by archive package description.
   - By name:
     i. Click the **Find** icon on the tool bar. The Find Files window opens.
     ii. Enter your search information in the Find Files window.
     iii. Click **Search**. The Matching Files window opens.
3. Click the selection boxes next to the objects that you want to retrieve.
4. Click **Retrieve**. The Retrieve Destination window opens.
5. Enter the information in the Retrieve Destination window and click **Retrieve**. The Retrieve Report window displays the processing results.

4.3 IBM Tivoli Storage Manager server for AIX

In this section we discuss the important steps for installation and configuration for AIX operating systems. We recommend that you read the IBM Tivoli Storage Manager quick start publications for the requirements for other UNIX flavors at the following URL:

http://publib.boulder.ibm.com/infocenter/tivihelp/index.jsp?toc=/com.ibm.itstorage.doc/toc.xml

The AIX server has the following requirements:

- An appropriately configured 64-bit POWER™ system computer
- AIX V5.3 or later running in a 64-bit kernel environment
- At least 200 MB of free disk storage
- Memory of 1 GB

This section is divided into the following sections:

- 4.3.1, “Tivoli Storage Manager installation packages for AIX” on page 162
- 4.3.2, “Installing Tivoli Storage Manager in an AIX environment” on page 165
- 4.3.3, “Installing IBM Tivoli Storage Manager Client for AIX” on page 169
- 4.3.4, “Installation verification” on page 171

4.3.1 Tivoli Storage Manager installation packages for AIX

You must stop all active server and storage agent processes before installation can proceed.Installation is stopped and an error message is displayed if an active server or storage agent process is detected.

Tivoli Storage Manager uses two types of AIX installation packaging. One AIX installation packaging is for a new release, and the other is for a maintenance update. It is important to understand the consequences of each type of installation.

- New install and migrate

  For a new version or release of Tivoli Storage Manager, install packages are set to perform a new install if no prior version is installed or a migrate install if a prior version is installed. A new install or a migrate install always establishes a base level for each file set.

  In the case of a migrate install, the prior base level and all updates to that base level are removed from the system. You cannot remove file sets installed by a migrate install without uninstalling the product.

- Maintenance update

  You can use an update installation package with a maintenance update (fix pack) to correct a small number of high-impact problems. An update installation allows the package to be rejected after installation and does not force a rebasing of a file set. Installation of a maintenance update uses the following options:

ALL Available Software	Install and Update
COMMIT software updates?	no
SAVE replaced files?	yes

Note: If you COMMIT a file set, you have just rebased this file set and cannot remove it without uninstalling the product.

Migrating to IBM Tivoli Storage Manager Version 5.5

It is possible to install Tivoli Storage Manager V5.5 over a previous version of ADSM or Tivoli Storage Manager. This process is called a *migrate install*. A DSMSERV UPGRADEDB operation is automatically performed during a migrate install.

Note: In a scenario with several installations of the server with a separate log and database, the upgrade database that is performed during the install has no effect. Enter the UPGRADEDB command for each database.

The default installation directories have changed for Tivoli Storage Manager. If you have previously used Disaster Recovery Manager to create a disaster recovery plan (DRP) file, this file refers to path names that are no longer valid. After you have installed Tivoli Storage Manager, you must back up your storage pools and database and create a new DRP file. For the sequence and details of the procedure, see the Disaster Recovery Manager chapter in the *Tivoli Storage Manager Administrator's Guide for AIX Administrator's Guide Version 5.5*, SC32-0117.

If you are using High-Availability Cluster Multi-Processing (HACMP) and you need to migrate to Tivoli Storage Manager V5.5 (from any previous version), you must convert to the new **startserver** script. Device resets are no longer performed by the **startserver** script. They are now performed by the server during the initialization of the library.

If you set the RESETDRIVES parameter to YES for a library, the reset is performed on the library manager for the library and all drives defined to it. If you define a library as SHARED, the RESETDRIVES parameter automatically defaults to YES for the library. Otherwise, you can run the UPDATE LIBRARY command with RESETDRIVES=YES.

To return to an earlier version of Tivoli Storage Manager after a migrate install, a full database backup of this original version is required along with the install code for the server of that original version. It is not possible to restore a later version's backed-up database onto the earlier version of the server.

In some cases, you might be required to migrate a previous version of the Tivoli Storage Manager server onto a newly built environment, on a separate machine with a higher version of Tivoli Storage Manager. We recommend that the new environment, where possible, be installed first with the same version of Tivoli Storage Manager server code, then restore the Tivoli Storage Manager database using the same version prior to the upgrade. This process ensures that the Tivoli Storage Manager database is upgraded accordingly.

Although it is possible for you to restore a Version 5.4 database onto a Tivoli Storage Manager V5.5 server, we do not recommend you do so because of the possible implications due to different database structures between the two versions.

Installation directories

In this section we present an overview of the following installation directories:

- Server directory (/usr/tivoli/tsm/server/bin), which contains:
  - Storage pool volumes (backup.dsm, archive.dsm, spcmgmt.dsm)
  - Database volume (db.dsm)

  - Recovery log volume (log.dsm)
  - Server code and licensing
- Related server directories:
  - The message repository, online help, and supported languages (/usr/lib/nls/msg/). The message repository is dsmserv.cat, and the command and message help is dsmserv.hlp
  - Tivoli inventory (/usr/tivoli/tsm/tivinv)
  - Event receiver (/etc/tivready/monitorslfs)
- Device directories:
  - /usr/tivoli/tsm/devices/bin
  - /etc/drivers
  - /etc/methods
  - /usr/tivoli/tsm/devices/bin and /etc)
- Language directory:

  For language-dependent portions of the program: /usr/lib/nls/msg/[lang]/ file. U.S. English, German, French, Italian, Spanish, Brazilian Portuguese, Korean, Japanese, traditional Chinese, simplified Chinese, Chinese GBK, Chinese Big5, Russian, Czechoslovakian, Hungarian, and Polish are supported. United States English (ISO8559), the en_US Cultural convention and Language translation environment, are required installations regardless of the language environment you use.

Installation restrictions

This section lists restrictions that you must consider prior to installing Tivoli Storage Manager. If an active server or storage agent process is detected, the installation stops, and an error message is displayed. All active server or storage agent processes must be halted before installation can proceed.

Installation of the Tivoli Storage Manager server and Tivoli Storage Manager storage agent on the same machine is not supported. During server installation, rmitab autosrvr is run, which can remove customized autosrvr entries. The installation can remove all lines with the identifier autosrvr even if you have customized other options and commands under this identifier.

4.3.2 Installing Tivoli Storage Manager in an AIX environment

Use the following procedure to install the server and device support.

1. Log on as the root user.
2. Load the CD into the selected drive.
   a. If you downloaded the code from the PartnerWorld or PassportAdvantage Web site, you can unzip the code by entering the following commands:

```
gunzip code.tar.gz
tar -xvf code.tar
```

   b. Change the directory to ../usr/sys/inst.images/.
   c. You might be required to create the installation .toc file using the **inutoc** command.

> **Note:** If you select a language other than U.S. English, see Appendix A, "National Language Support" in the *IBM Tivoli Storage Manager for AIX Installation Guide Version 5.5*, GC23-5969, for the correct selection path:
>
> http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/topic/com.ibm.itsmfdt.doc/b_install_guide_aix53.htm#r_nls

3. Run the command **smitty**, and select **Software Installation and Maintenance → Install and Update Software → Install and Update from ALL Available Software**.
4. Select the INPUT device that you are using for the installation. You can enter the drive name in the window or press F4 to access the device list and Select the device / directory for software.
5. Locate the line **Software to Install**. Press PF4. The value of the environment variable LANG controls which translated messages, help, and device-selection messages in SMIT are automatically installed.

6. Select the file sets that you want to install for your configuration. Press PF7, and then click **OK**. Table 4-4 lists the current file sets available.

> **Note:** Take into consideration the following recommendations:
>
> - If you are not sure that your system currently has the Version 8 C++ Runtime, select the Version 8 C++ Runtime file set (xlC.rte.8.0.0.0).
> - If you are not sure that your system currently has the =GSKit Version 7.0.4.11 file set (gsksa.rte).

Table 4-4 AIX file set list

File set name	Description
tivoli.tsm.devices.acls	STK Silo support (optional)
tivoli.tsm.devices.aix5.rte	Device support for AIX V5.3 or later, 64 bit
tivoli.tsm.license.aix5.rte64	License enablement module, 64 bit kernel
tivoli.tsm.license.cert	License certificates
tivoli.tsm.loc.server.ela	Electronic license agreements
tivoli.tsm.msg.en_US.devices	SMIT menu catalogs, U.S. English
tivoli.tsm.msg.en_US.server	Message library and help
tivoli.tsm.server.aix5.rte64	Server runtime, AIX V5.3 or later, 64 bit
tivoli.tsm.server.com	Server samples and diagnostic utilities
tivoli.tsm.server.webcon	Web console
xlC.rte.8.0.0.0	C++ Runtime for AIX Version 8
xlC.msg.en_US.rte.8.0.0.0	Messages for C++ Runtime for AIX Version 8
gsksa.rte	64-bit server encryption and SSL capabilities

7. You must accept new license agreements to install the server. To view the license agreement, select yes on the Preview new LICENSE agreements? option. To accept the license agreement, select yes on the Accept new LICENSE agreements? option. The license agreements are translated in many languages. The language environment that you set determines the version displayed (see Figure 4-26).

```
                 Install and Update from ALL Available Software

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                     [Entry Fields]
* INPUT device / directory for software               .
* SOFTWARE to install                                [gksa                  > +
  PREVIEW only? (install operation will NOT occur)    no                      +
  COMMIT software updates?                            yes                     +
  SAVE replaced files?                                no                      +
  AUTOMATICALLY install requisite software?           yes                     +
  EXTEND file systems if space needed?                yes                     +
  OVERWRITE same or newer versions?                   no                      +
  VERIFY install and check file sizes?                no                      +
  DETAILED output?                                    no                      +
  Process multiple volumes?                           yes                     +
  ACCEPT new license agreements?                      yes                     +
  Preview new LICENSE agreements?                     no                      +


F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 4-26 Accept new license agreement

8. Check the default settings for the options on this window. To continue, press Enter. Confirm the message, and press Enter again. This installation can take several minutes. The installation status is shown in Figure 4-27.

```
                                COMMAND STATUS

Command: OK            stdout: yes           stderr: no

Before command completion, additional instructions may appear below.

[TOP]
geninstall -I "a -cgNqwXY -J"  -Z   -d . -f File 2>&1

File:
    I:gsksa.rte                   7.0.4.11
    I:tivoli.tsm.devices.acsls     5.5.0.0
    I:tivoli.tsm.devices.aix5.rte  5.5.0.0
    I:tivoli.tsm.license.aix5.rte64 5.5.0.0
    I:tivoli.tsm.license.cert      5.5.0.0
    I:tivoli.tsm.loc.server.ela    5.5.0.0
    I:tivoli.tsm.msg.EN_US.devices 5.5.0.0
    I:tivoli.tsm.msg.EN_US.server  5.5.0.0
    I:tivoli.tsm.msg.en_US.devices 5.5.0.0
[MORE...398]

F1=Help             F2=Refresh          F3=Cancel           F6=Command
F8=Image            F9=Shell            F10=Exit            /=Find
n=Find Next
```

Figure 4-27 Installation status

> **Note:** Some file sets are OS level-specific and might fail to install. This is normal behavior, so do not interpret it as an error.

9. After the installation is complete, exit the SMIT. The database, Recovery log, and default storage pools are created by the installation process.

To start the server from the /usr/tivoli/tsm/server/bin directory, run the command **`nohup ./dsmserv`**. It is also required to install a backup/archive client and an administrative client.

When you install or upgrade the server to Tivoli Storage Manager V5.5, you must give your servers unique names. If all of your servers have the default name, `SERVER1`, you can only add one of them to the Administration Center. Use the SET SERVERNAME command to specify the new server name. The maximum length of the name is 64 characters.

4.3.3 Installing IBM Tivoli Storage Manager Client for AIX

If you are migrating from prior versions of the client, do a migrate install, or uninstall the older client and then install the new client. The client can be used on either 32-bit or 64-bit machines, separate installation packages are no longer available.

Due to the compiler upgrade for the AIX client from Tivoli Storage Manager V5.4, the newer version of libC (CSet++) runtime is required. The Tivoli Storage Manager AIX client package includes all of the required packages, and they overwrite the older runtime on your system during installation.

For detailed information about how to install the Tivoli Storage Manager V5.5 backup/archive client, see *Tivoli Storage Manager for UNIX and Linux: Backup-Archive Clients Installation and User Guide*, SC32-0145-01, available at:

`http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/topic/com.ibm.itsmfdt.doc/ans50000.htm`

To install Tivoli Storage Manager from a CD, complete the following steps:

1. Log in as the root user. Insert the CD into the CD drive device. You must manually mount the CD device for remote or NFS installations.
2. From the AIX command line, type `smitty install` and press Enter.
3. Select **Install and Update Software**, and press Enter.
4. Select **Install and Update From ALL Available Software**, and press Enter.
5. At the INPUT device/directory for software prompt, press the F4 key and select the CD device containing the installation CD, or specify the directory containing the installation images, and press Enter.

6. At the SOFTWARE to install prompt, press the F4 key. Select the Tivoli Storage Manager file sets that you want to install by pressing the F7 key and press Enter (see Figure 4-28).

```
                 Install and Update from ALL Available Software

Ty+--------------------------------------------------------------------------+
Pr|                          SOFTWARE to install                             |
  |                                                                          |
  | Move cursor to desired item and press F7. Use arrow keys to scroll.      |
* |     ONE OR MORE items can be selected.                                   |
* | Press Enter AFTER making all selections.                                 |
  |                                                                          |
  | [MORE...19]                                                              |
  | > tivoli.tsm.client.ba                                               ALL |
  |     + 5.5.0.0  TSM Client - Backup/Archive Base Files                    |
  |     + 5.5.0.0  TSM Client - Backup/Archive Common Files                  |
  |     + 5.5.0.0  TSM Client - Backup/Archive WEB Client                    |
  |     + 5.5.0.0  TSM Client - IMAGE Backup Client                          |
  |     + 5.5.0.0  TSM Client - NAS Backup Client                            |
  |                                                                          |
  | > tivoli.tsm.client.ba.gpfs                                          ALL |
  | [MORE...84]                                                              |
  |                                                                          |
  | F1=Help                 F2=Refresh              F3=Cancel                |
F1| F7=Select               F8=Image                F10=Exit                 |
F5| Enter=Do                /=Find                  n=Find Next              |
F9+--------------------------------------------------------------------------+
```

Figure 4-28 List Tivoli Storage Manager client file sets to install

7. Select the options you want, and press Enter to begin the installation. After installation completes, you see a completion panel (see Figure 4-29).

```
                                COMMAND STATUS

Command: OK            stdout: yes           stderr: no

Before command completion, additional instructions may appear below.

[TOP]
geninstall -I "a -cgNqwXY -J"  -Z   -d . -f File 2>&1

File:
    I:gsksa.rte                    7.0.4.11
    I:tivoli.tsm.client.api.64bit  5.5.0.0
    I:tivoli.tsm.client.ba.32bit.base 5.5.0.0
    I:tivoli.tsm.client.ba.32bit.common 5.5.0.0
    I:tivoli.tsm.client.ba.32bit.image 5.5.0.0
    I:tivoli.tsm.client.ba.32bit.nas 5.5.0.0
    I:tivoli.tsm.client.ba.32bit.web 5.5.0.0
    I:tivoli.tsm.client.msg.JA_JP  5.5.0.0
    I:tivoli.tsm.client.msg.Ja_JP  5.5.0.0
[MORE...355]

F1=Help             F2=Refresh          F3=Cancel           F6=Command
F8=Image            F9=Shell            F10=Exit            /=Find
n=Find Next
```

Figure 4-29 Client installation status

4.3.4 Installation verification

The installation verification performs some common functions with the Tivoli Storage Manager server and client. We discuss the functions in the following sections:

- "Starting and stopping the server" on page 171
- "Registering licenses" on page 172
- "Query commands" on page 172
- "Configuring client" on page 172
- "Backing up files from the client" on page 173

Starting and stopping the server

To start the server from the /usr/tivoli/tsm/server/bin directory, enter `nohup ./dsmserv` (or use `nohup ./dsmserv &` to run in the background).

You can stop the server without warning if an unexpected problem requires you to return control to the operating system. To avoid losing administrative and client node connections, stop the server only after current sessions have completed or been canceled.

To stop the server, issue the **halt** command from the Tivoli Storage Manager command prompt.

Registering licenses

You must immediately register any Tivoli Storage Manager licensed functions that you purchase. Use the REGISTER LICENSE command to do so. Also, be sure to read the Tivoli Storage Manager readme file for the latest information regarding enrollment certificate files.

Query commands

It is possible to verify server installation by issuing the following QUERY commands from the server console:

- QUERY STATUS
- QUERY STGPOOL: Shows the stgpools that are automatically set up by the installation
- QUERY DB: Shows a basic database that is set up by the installation
- QUERY LOG: Shows a basic log that is set up by the installation
- QUERY OPTION: Shows options that were set up in the dsmserv.opt file

After the server is installed, it runs in the background, so you must start the server in the foreground or use an administrative client to issue the commands.

Backing up client data from a backup/archive client can be done to verify the complete Tivoli Storage Manager installation. The client must have been installed from the UNIX client CD.

Configuring client

The server must be running before you can use the administrative and backup/archive clients. Configure the backup/archive client by completing the following steps:

1. Copy the sample client system options file (dsm.sys.smp) and the sample client user options file (dsm.opt.smp). The sample files are in /usr/tivoli/tsm/client/ba/bin/.
2. Edit the options files to include the following options:
   - In dsm.opt:

     ```
     servername server_name
     ```

   - In dsm.sys:

     ```
     servername server_name
     commmethod tcpip
     ```

```
tcpport port_address
tcpserveraddress server_address
nodename client
```

Note: The server names specified in dsm.opt and dsm.sys must match.

3. Start the backup/archive client graphical user interface by entering `dsm`.

   The default ID and password for the backup/archive client are both `client`.

Backing up files from the client

To back up several files, perform the steps described in this section.

Note: Do not select files, which are to be backed up, that exceed your storage pool capacity. The backup storage pool, archive storage pool, and space management pool are allocated for 8 MB during installation.

1. Click **Backup**. The Backup window is displayed. You can select files from a directory tree, or you search a drive or directory to select files.
   - To select files from a directory tree, expand the tree and select the files you want to back up.
   - To search for files:
     i. Click the **Find** icon on the tool bar.
     ii. The Find Files (Backup) window is displayed. Enter your search criteria.
     iii. Click **Search**.
     iv. The Matching Files (Backup) window is displayed. Click the selection boxes next to the files that you want to back up.
     v. Close the Matching Files (Backup) window by clicking the **Backup/Archive** icon on the menu bar and selecting **Close**.
2. Click **Always Backup** from the Type of Backup list.
3. Click **Backup**. The Backup Status window displays the backup processing status. A message is displayed when the backup has completed successfully.

4.4 IBM Tivoli Storage Manager for Linux Server

For references about the requirements for installing Tivoli Storage Manager on Linux, refer to 2.1.1, “IBM Tivoli Storage Manager supported server OS and hardware” on page 40. This section discusses the installation of the Tivoli Storage Manager on a Linux platform. The discussion is divided into the following sections:

4.4.1 Installation directories in Linux

The installation of a Tivoli Storage Manager server in Linux uses the following installation directories:

- Default server directory (/opt/tivoli/tsm/server/bin), which contains:
  - Storage pool volumes (backup.dsm, archive.dsm, spcmgmt.dsm)
  - Database volume (db.dsm)
  - Recovery log volume (log.dsm)
  - Server and licensing
- Related server directories:
  - The message repository, online help, and supported languages (/opt/tivoli/tsm/server/bin/en_US)
  - Device support (/opt/tivoli/tsm/devices/bin)
  - IBM Tivoli inventory (/opt/tivoli/tsm/tivinv)

4.4.2 Installing Tivoli Storage Manager in a Linux environment

Installing and configuring the Tivoli Storage Manager server and device support requires that you follow a sequence of steps. To install the server on Linux, you must accept the license agreement. Perform the following steps:

1. Log on as the root user.
2. Navigate to the directory where the Tivoli Storage Manager server package was uncompressed and unpacked or to the CDROM directory.

3. To install the server on Linux, you must accept the license agreement. Perform the following steps.

   a. Run the install_server script. The syntax is:

   ```
   ./install_server [-p prefix][platform][-j <location of JRE>]
   ```

   where:

   - `prefix` is optional, and it denotes the file system where you want to install Tivoli Storage Manager. If you do not provide a prefix, the default is /opt.
   - `platform` is optional and denotes the Tivoli Storage Manager architecture to be installed (ppc64, i686, s390x, x86_64, ia64).
   - Location of JRE™ is optional and denotes the location of a JRE on your system. You must use it only if the bundled JRE fails to display the license agreement.

   b. The **`install_server`** script displays the license agreement in the language set by your operating system and prompts you to accept it. To accept the license agreement, type `1` followed by pressing Enter.

4. After you have accepted the license agreement, a menu is displayed with all the available installation packages for the corresponding platform. You can quit and install the package manually by typing `q` followed by pressing Enter, or you can run the basic installation by typing `b` followed by pressing Enter.

   – Manual installation can use the **rpm** command, and using either the `-i` option for a new installation or the `-U` option for upgrade install. You can also change the installation path prefix using the --prefix parameter. The default path is /opt. You can also remove an existing package using the `-e` option. Checking existing packages can be performed using the following command:

   ```
   rpm -qa | grep TIVsm.
   ```

   – The basic installation detects any existing installation and automatically chooses to use the **`rpm -i`** or **`rpm -U`** command.

5. One directory contains the rpm packages for each architecture supported by Tivoli Storage Manager on Linux. The directory paths are:

ia64	Packages for Itanium Red Hat Enterprise Linux 4 and 5; SUSE Linux Enterprise Server 9 and 10
i686	Packages for Intel-based Red Hat Enterprise Linux 4 and 5; SUSE Linux Enterprise Server 9 and 10
s390x	Packages for System z-based Red Hat Enterprise Linux 4 and 5; SUSE Linux Enterprise Server 9 and 10

ppc64/glibc-2.3 — Packages for Power systems-based Red Hat Enterprise Linux 4; SUSE Linux Enterprise Server 9 and 10

noarch — Packages for supporting various languages other than English; packages to be installed on any Tivoli Storage Manager supported Linux environment

> **Note:** Installation of the Tivoli Storage Manager server and Tivoli Storage Manager Storage Agent on the same machine is not supported.

You must install the Web administrator, license, and Web help packages under the same file system as the Tivoli Storage Manager server. The default path is /opt/tivoli/tsm/server/bin. The packages are listed in Table 4-5.

Table 4-5 Iinstall packages

	Package	Description
IA32	TIVsm-server-5.5.0-0.i386.rpm	Server package for IA32
	TIVsm-license-5.5.0-0.i386.rpm	License enabler for IA32
	TIVsm–tsmscsi–5.5.0-0-i386.rpm	SCSI pass-through drivers and utilities
IA64	TIVsm-server-5.5.0-0.ia64.rpm	Server package for IA64
	TIVsm-license-5.5.0-0.ia64.rpm	License enabler for IA64
	TIVsm–tsmscsi–5.5.0-0-ia64.rpm	SCSI pass-through drivers and utilities
ppc	TIVsm-server-5.5.0-0.ppc64.rpm	Server package for ppc64
	TIVsm-license-5.5.0-0.ppc64.rpm	License enabler for ppc64
	TIVsm-tsmscsi-5.5.0-0-ppc64.rpm	SCSI pass-through drivers and utilities
s390x	TIVsm-server-5.5.0-0.s390x.rpm	Server package for s/390x
	TIVsm-license-5.5.0-0.s390x.rpm	License enabler for s/390x
	TIVsm-tsmscsi-5.5.0-0-s390x.rpm	SCSI pass-through drivers and utilities

6. When you install or upgrade the server to Tivoli Storage Manager V5.5, you must give your servers unique names. If all of your servers have the default name, SERVER1, you are restricted to adding only one of them to the Administration Center. Use the SET SERVERNAME command to specify the

new server name. The maximum length of the name is 64 characters. To name the server ITSO_SERVER, use this command: **`set servername ITSO_SERVER`**. See the *IBM Tivoli Storage Manager for Linux Administrator's Reference Version 5.5*, SC32-0125, for more details about this command.

7. When the installation is finished, remove the CD from the drive.

> **Note:** Messages and help install to /opt/tivoli/tsm/server/bin/<*locale*>, where <*locale*> is one of the locales specified. The message catalogs are always called dsmserv.cat, and the help files are called dsm<*locale*>.hlp.

Table 4-6 lists Tivoli Storage Manager messages and command-line help. Use these packages to customize your installation after the basic installation is complete. These packages are architecture independent. They can be installed on any Linux platforms.

Table 4-6 Tivoli Storage Manager messages and command-line help

Package name	Components
TIVsm-cmdlinehelp.de_DE-5.5.0.0.noarch.rpm	German
TIVsm-cmdlinehelp.es_ES-5.5.0.0.noarch.rpm	Spanish
TIVsm-cmdlinehelp.de_FR-5.5.0.0.noarch.rpm	French
TIVsm-cmdlinehelp.it_IT-5.5.0.0.noarch.rpm	Italian
TIVsm-cmdlinehelp.ja_JP-5.5.0.0.noarch.rpm	Japanese
TIVsm-cmdlinehelp.ko_KR-5.5.0.0.noarch.rpm	Korean
TIVsm-cmdlinehelp.pt_BR-5.5.0.0.noarch.rpm	Brazilian Portuguese
TIVsm-cmdlinehelp.zh_TW-5.5.0.0.noarch.rpm	Traditional Chinese (big5 codepage)
TIVsm-cmdlinehelp.zh_TW.euctw-5.5.0.0.noarch.rpm	Traditional Chinese (euctw codepage)
TIVsm-cmdlinehelp.zh_CN-5.5.0.0.noarch.rpm	Simplified Chinese (eucCN codepage)
TIVsm-cmdlinehelp.zh_CN.gb18030-5.5.0.0.noarch.rpm	Simplified Chinese (gb18030 codepage)
TIVsm-cmdlinehelp.de_DE.utf8-5.5.0-0	German uft8
TIVsm-cmdlinehelp.es_ES.utf8-5.5.0-0	Spanish uft8
TIVsm-cmdlinehelp.fr_FR.utf8-5.5.0-0	French uft8
TIVsm-cmdlinehelp.it_IT.utf8-5.5.0-0	Italian uft8
TIVsm-cmdlinehelp.ja_JP.utf8-5.5.0-0	Japanese uft8

Package name	Components
TIVsm-cmdlinehelp.ko_KR.utf8-5.5.0-0	Korean uft8
TIVsm-cmdlinehelp.pt_BR.utf8-5.5.0-0	Brazilian Portuguese uft8
TIVsm-cmdlinehelp.zh_TW.utf8-5.5.0-0	Traditional Chinese uft8
TIVsm-cmdlinehelp.zh_CN.utf8-5.5.0-0	Simplified Chinese uft8
TIVsm-cmdlinehelp.en_US.utf8-5.5.0-0	United States English uft8

To ensure compatibility between the server and the storage agent, check the Tivoli Web site at:

http://www-1.ibm.com/support/docview.wss?uid=swg21302789

This site has version requirements. The Tivoli Storage Manager servers for the z/OS, AIX, HP-UX, Linux, Sun Solaris, and Microsoft Windows operating systems support LAN-free data movement.

When upgrading multiple servers participating in library sharing to Version 5.5 and later, the servers acting as library manager must be upgraded first to maintain compatibility among the servers acting as library clients or storage agents. When upgrading multiple servers participating in server-to-server connection to Version 5.5 and later, all servers must be upgraded at the same time.

The Tivoli Storage Manager server supports storage agents and library clients at the same version, release, and modification level, including the previous version.

Note: An exception to this rule is where a fix or a product enhancement requires concurrent code changes to server, storage agent, and library client.

Version 5.5 servers acting as a library manager are compatible with Version 5.4 servers acting as a library client. However, Version 5.5 servers acting as library clients remain incompatible with Version 5.4 servers acting as library managers.

4.4.3 Installation verification

Starting and stopping the server

To start the server from the /usr/tivoli/tsm/server/bin directory, enter the following command to run in the background:

```
nohup ./dsmserv (or use nohup ./dsmserv &
```

You can stop the server without warning if an unexpected problem requires you to return control to the OS. To avoid losing administrative and client node connections, stop the server only after current sessions have completed or been canceled.

To stop the server, issue the `halt` command from the Tivoli Storage Manager command prompt.

Registering licenses

You must immediately register any Tivoli Storage Manager licensed functions that you purchase. Use the REGISTER LICENSE command to do this, as documented in *IBM Tivoli Storage Manager for Linux Administrator's Reference Version 5.5*, SC32-0125, and *IBM Tivoli Storage Manager for Linux Administrator's Guide Version 5.5*, SC32-0119. Also, be sure to refer to the Tivoli Storage Manager readme file for the latest information regarding enrollment certificate files.

Using the query commands

It is possible to verify server installation by issuing QUERY commands from the server console. After being installed, the server runs in the background, so you must start the server in the foreground or use an administrative client to issue the commands. You can issue the following QUERY commands:

- QUERY STATUS
- QUERY STGPOOL: Shows the stgpools that the installation automatically sets up
- QUERY DB: Shows a basic database that the installation sets up
- QUERY LOG: Shows a basic log that the installation sets up

Performing backup

You can back up client data from a backup/archive client to verify the complete Tivoli Storage Manager installation. The client must have been installed from the UNIX client CD.

> **Note:** Do not select files, which are to be backed up, that exceed your storage pool capacity. An 8 MB backup storage pool, an 8 MB archive storage pool, and an 8 MB space management pool were allocated at installation.

The server must be running before you can use the administrative and backup/archive clients. To configure the backup/archive client, complete the following steps:

1. Copy the sample client system options file (dsm.sys.smp) and the sample client user options file (dsm.opt.smp). The sample files are in /usr/tivoli/tsm/client/ba/bin/.
2. Edit the options files to include the following options:
   – In dsm.opt:

     ```
     servername server_name
     ```

   – In dsm.sys:

     ```
     servername server_name
     commmethod tcpip
     tcpport port_address
     tcpserveraddress server_address
     nodename client
     ```

     Note: The server names specified in dsm.opt and dsm.sys must match.

3. Start the backup/archive client graphical user interface by entering the **dsm** command.

   The default ID and password for the backup/archive client are both client.

To back up several files, perform the following steps:

1. Click **Backup**. The Backup window is displayed.

   You can select files from a directory tree or search a drive or directory to select files:

   a. To select files from a directory tree:

      Expand the tree and select the files you want to back up.

   b. To search for files:

      i. Click the **Find** icon on the tool bar.

      ii. The Find Files (Backup) window opens. Enter your search criteria in the Find Files (Backup) window and click **Search**.

      iii. The Matching Files (Backup) window is displayed. Click the selection boxes next to the files that you want to back up.

      iv. Close the Matching Files (Backup) window by clicking the **Backup/Archive** icon on the menu bar and selecting **Close**.

2. Click **Always Backup** from the Type of Backup list.
3. Click **Backup**. The Backup Status window displays the backup processing status. A message displays when the backup has completed successfully.

4.5 Getting started with administrative tasks

This section provides an introduction to some basic IBM Tivoli Storage Manager administrative tasks and discusses the following topics:

- 4.5.1, "Managing the Tivoli Storage Manager server" on page 181
  - "Starting the Tivoli Storage Manager server" on page 181
  - "Stopping the Tivoli Storage Manager server" on page 182
  - "Backing up the server database and database recovery log" on page 183
- 4.5.2, "Managing Tivoli Storage Manager client schedules" on page 183
  - "Starting the Tivoli Storage Manager scheduler" on page 183
  - "Verifying a schedule" on page 183
- 4.5.3, "Setting the client/server communications options" on page 183
- 4.5.4, "Managing Tivoli Storage Manager administrators" on page 185
  - "Registering additional administrators" on page 185
  - "Changing administrator passwords" on page 186
- 4.5.5, "Returning to a previous version of Tivoli Storage Manager" on page 186

You can also use the Administration Center to manage servers and clients.

4.5.1 Managing the Tivoli Storage Manager server

This section describes managing functions for the Tivoli Storage Manager server.

Starting the Tivoli Storage Manager server

You can start the Tivoli Storage Manager server in several ways. However, we recommend that you start it as a service. In this way, the server remains active when you log off the workstation. To start the server as a service, perform these steps from the Tivoli Storage Manager Console:

1. Expand the tree for the Tivoli Storage Manager server that you are starting and expand **Reports**.
2. Click **Service Information**. The Service Information view is displayed in the right pane.

3. If the server status displays Stopped, right-click the service line and select **Start**.

Stopping the Tivoli Storage Manager server

You can stop the server without warning if required. To avoid losing administrative and client node connections, stop the server only after current sessions have been completed or canceled.

> **Note:** This procedure shuts down the server immediately. The shutdown also cancels all client sessions. All other procedures in this book assume that the server is running. Shutting down the server is only performed for maintenance purposes.

To stop the server, perform the following steps:

- Stop a server that is running as a service:
  a. Expand the tree for the Tivoli Storage Manager server you are stopping, and expand **Reports**.
  b. Click **Service Information**. The Service Information view is displayed in the right pane.
  c. Right-click the server service line and select **Stop**.
- Stop a server from the administrative Web interface:
  a. From the tree view in the browser, expand **Object View**.
  b. Expand **Server**.
  c. Click **Server Status**.
  d. From the pull-down menu, select **Halt Server** and click **Finish**.
- Stop a server from the administrative command line:
  a. Expand the tree for the Tivoli Storage Manager server that you are stopping and expand **Reports**.
  b. Click **Command Line**. The Command Line view is displayed in the right pane.
  c. Click **Command Line Prompt** in the right pane. The Command Prompt dialog is displayed.
  d. Enter `halt` in the Command field, and click **Submit**.

Backing up the server database and database recovery log

If the Tivoli Storage Manager server database or the recovery log is unusable, the entire server is unavailable. If a database is lost and cannot be recovered, all of the data managed by this server is lost. If a storage pool volume is lost and cannot be recovered, the data on the volume is also lost.

With Tivoli Storage Manager, you can define administrative schedules so that the database and storage pools are backed up regularly. If you lose your database or storage pool volumes, you can use offline utilities provided by Tivoli Storage Manager to restore your server and data.

4.5.2 Managing Tivoli Storage Manager client schedules

This section shows how to start Tivoli Storage Manager schedules that you have defined and how to verify that they are running correctly.

Starting the Tivoli Storage Manager scheduler

The Tivoli Storage Manager client scheduler is the client component of the Tivoli Storage Manager scheduling model. The client scheduler runs as a Microsoft Windows service and must be installed and running on the Tivoli Storage Manager client machine to execute any client schedules you define to the Tivoli Storage Manager server. You can install the client scheduler using a wizard provided by the Tivoli Storage Manager client graphical interface. You can manually start the Scheduler service on each client node, or update the managed services option in the client options file to automatically start the Scheduler service as required. Refer to the *Backup-Archive Client Installation and User's Guide* that corresponds to your operating system for more information.

Verifying a schedule

You can verify that automation is working as expected on the day after you define the schedule and associate it with clients. If the schedule runs successfully, the status shows it as Completed.

> **Note:** The include-exclude list (file on UNIX clients) on each client also affects which files are backed up or archived.

4.5.3 Setting the client/server communications options

Use the Server Options utility on the Tivoli Storage Manager console to view and specify the server communications options. This utility is available from the Service Information view in the server tree. By default, the server uses the

TCP/IP, named pipes, and HTTP communication methods. If you start the server console and see warning messages that indicate a protocol cannot be used by the server, either the protocol is not installed or the settings do not match the Microsoft Windows protocol settings.

For a client to use a protocol that is enabled on the server, the client options file must contain corresponding values for communication options. From the Server Options utility, you can view the values for each protocol.

> **Attention:** This section describes setting server options before you start the server. When you start the server, the new options go into effect. If you modify any server options after starting the server, you must stop and restart the server to activate the updated options.

For more information about server options, see the *IBM Tivoli Storage Manager for Linux Administrator's Reference Version 5.5*, SC32-0125, or the Tivoli Storage Manager console online help.

TCP/IP options

Example 4-2 shows an example of a TCP/IP setting.

Example 4-2 TCP/IP setting

```
commmethod tcpip
tcpport 1500
tcpwindowsize 63
tcpnodelay no
```

Named pipes options

The named pipes communication method is ideal when running the server and client on the same Microsoft Windows machine because named pipes support is internal to the Windows base system. Named pipes require no special configuration. Example 4-3 shows the named pipes settings.

Example 4-3 Named pipe settings

```
commmethod namedpipe
namedpipename \\.\pipe\adsmpipe
```

SNMP DPI subagent options

Tivoli Storage Manager implements a simple network management protocol (SNMP) subagent. You can configure the SNMP subagent to send traps to an SNMP manager, such as IBM Tivoli Network Manager, and to provide support for a management information base (MIB).

The subagent communicates with the snmpd daemon, which in turn communicates with a management application. The snmpd daemon must support the IBM DPI® protocol. Agents are available on AIX. The subagent process is separate from the Tivoli Storage Manager server process, but the subagent gets its information from a server options file. When the SNMP management application is enabled, it can get information and messages from servers.

Example 4-4 shows the SNMP setting. You must specify the COMMMETHOD option.

Example 4-4 SNMP settings

```
commmethod snmp
snmpheartbeatinterval 5
snmpmessagecategory severity
```

4.5.4 Managing Tivoli Storage Manager administrators

This section describes how to manage Tivoli Storage Manager administrators.

Registering additional administrators

If you are adding administrators, you must register them and grant an authority level to each.

> **Note:** The name SERVER_CONSOLE is reserved for Tivoli Storage Manager console operations and cannot be used as the name of an administrator.

From the administrative Web interface, perform the following steps to register an administrative client and grant an authority level:

1. From the tree view, expand **Administrators**.
2. From the Operations pull-down menu, select **Register an Administrator**.
3. Enter the required information and click **Finish**.

Changing administrator passwords

From the administrative Web interface, perform the following steps to change the password of an administrator:

1. From the tree view, expand **Administrators**.
2. Select an administrator name.
3. From the Operations pull-down menu, select **Update an Administrator**.
4. Enter the password and click **Finish**.

4.5.5 Returning to a previous version of Tivoli Storage Manager

This section explains how to return to a previous version of the Tivoli Storage Manager in case of some unforeseen issues with the upgrade.

We describe what you must do before you remove the current version of the Tivoli Storage Manager server and how to return to a previous version of Tivoli Storage Manager (or ADSM). After you remove the Tivoli Storage Manager device driver, you have to restart your system.

To return to ADSM or an earlier version of Tivoli Storage Manager after you perform a migrate install, perform a full database backup from your original version and save the server install code for your original version.

> **Note:** You cannot restore a backed-up database from a prior version onto a newer version of the Tivoli Storage Manager server.

If you return to ADSM or an earlier version of Tivoli Storage Manager, be aware of the following results:

- References to client files that were backed up, archived, or migrated to the current Tivoli Storage Manager server are lost.
- Some volumes might be overwritten or deleted during Tivoli Storage Manager server operation. If so, client files that were on these volumes and that were migrated, reclaimed, moved (MOVE DATA command), or deleted (DELETE VOLUME command) might not be accessible to the earlier version of ADSM or Tivoli Storage Manager.
- Definitions, updates, and deletions of Tivoli Storage Manager objects that were performed on the current Tivoli Storage Manager server are lost.

Before you remove Tivoli Storage Manager, follow these steps:

1. Perform a full database backup.

   If you have a tape device class named tapeclass, enter this command to perform a full backup:

   ```
   backup db type=full devclass=tapeclass
   ```

2. Save a copy of the volume history and device configuration files that you defined on the VOLHISTORY and DEVCONFIG options in the server options file.

   To save the volume history in a file named volhist, and the device configuration in a file named devices, enter:

   ```
   backup volumehistory filenames=volhist backup devconfig
   filenames=devices
   ```

3. Store the output volumes in a safe location.

4.6 Integrated Solutions Console

This section provides the step-by-step procedure for installing the Integrated Solution Console (ISC) and Administration Center. Refer to the appropriate Tivoli Storage Manager installation guide for detailed instructions on this process.

Before you install the Integrated Solutions Console and the Administration Center, you must install, configure, and start the Tivoli Storage Manager server. You must install Integrated Solutions Console Version 6.0.1.1 before installing the Administration Center.

The installation wizards for the Integrated Solutions Console and the Administration Center are on separate CDs. The Integrated Solutions Console is only available through your regular order system such as Passport Advantage or the Advanced Administrative System (AAS).

After both applications are installed, you are able to administer Tivoli Storage Manager from anywhere in your network through a Web browser.

You can obtain the Release Notes from the IBM Tivoli Storage Manager Information Center:

http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/topic/com.ibm.itsmreadme.doc_5.5/relnote_adminctr550.html

4.6.1 Implementing ISC and Administration Center

The administrative Web interface is replaced in this release with the Administration Center. The Administration Center is a Web-based interface that you can use to centrally configure and manage Tivoli Storage Manager servers. You can use the Administration Center only to administer Version 5.4 or later servers. The old administrative Web interface cannot be used with Version 5.4 or later servers.

The Administration Center is installed as an plugin to the IBM Integrated Solution Console. The ISC enables you to install components provided by multiple IBM applications and access them from a single interface.

We provide the following six basic steps for setting up the Integrated Solutions Console and the Administration Center:

1. Install and start your Tivoli Storage Manager Version 5.5 servers. Give each server a unique name.
2. Use the Integrated Solutions Console CD or image to install the ISC. During the installation process, create an Integrated Solutions Console user ID and password.
3. Use the Administration Center CD or image to install the Administration Center on the same system as the ISC.
4. Log onto the ISC using a Web browser. You can then administer Tivoli Storage Manager from a browser anywhere in your network.
5. Add connections for the Tivoli Storage Manager servers you want to manage.
6. The default Integrated Solutions Console ID is iscadmin. You must specify a password for this ID. You can create additional ISC user IDs and passwords for other administrators to access the Administration Center.

The Tivoli Storage Manager server can require a large amount of memory, network bandwidth, and processor resources. In many cases, the server performs best when other applications are not installed on the same system. If the system meets the combined requirements for the server and the Administration Center, it can support both applications.

If you plan to use the Administration Center to manage an environment with a large number of servers or administrators, consider installing the Administration Center on a separate system.

For Administration Center system requirements, see the following Web site:

http://www.ibm.com/support/docview.wss?uid=swg21195062

4.6.2 Installing Integrated Solutions Console

The ISC installation media and the Administration Center must be extracted to different directories.

To install the ISC from a CD, complete the following steps:

1. Run the **`setupISC`** command from the installation image directory.
2. After the installation preparation wizard completes, you see the welcome window, similar to the one shown in Figure 4-30. Click **Next**.

Figure 4-30 Welcome window

3. In the License Agreement window, read the license carefully, accept the agreement, and click **Next**.

4. Enter the user name you want to create for ISC admin console, and enter and confirm a password (see Figure 4-31). Click **Next**.

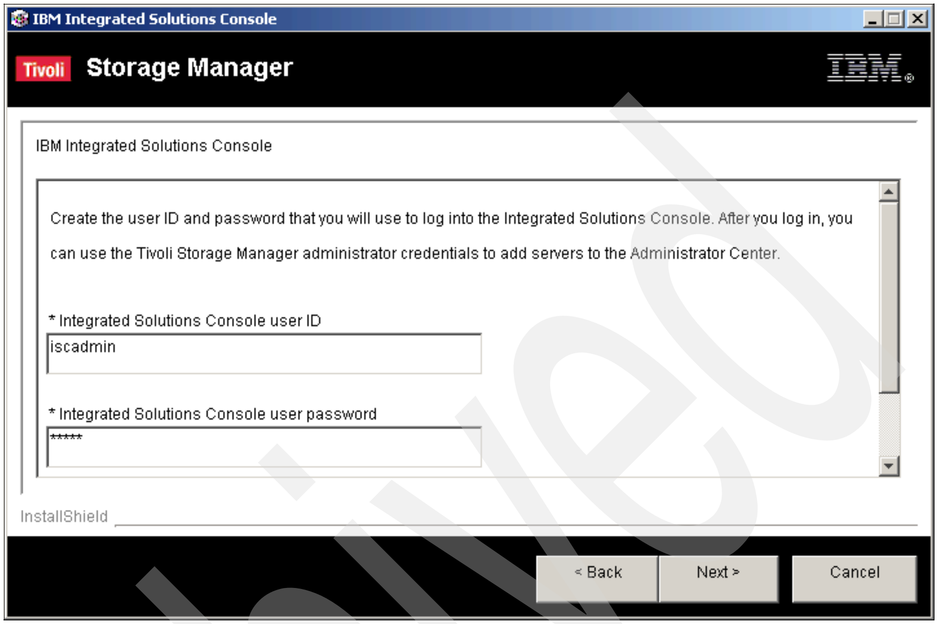

Figure 4-31 ISC admin user creation

5. Select the destination drive and folder, as shown in Figure 4-32, and click **Next**.

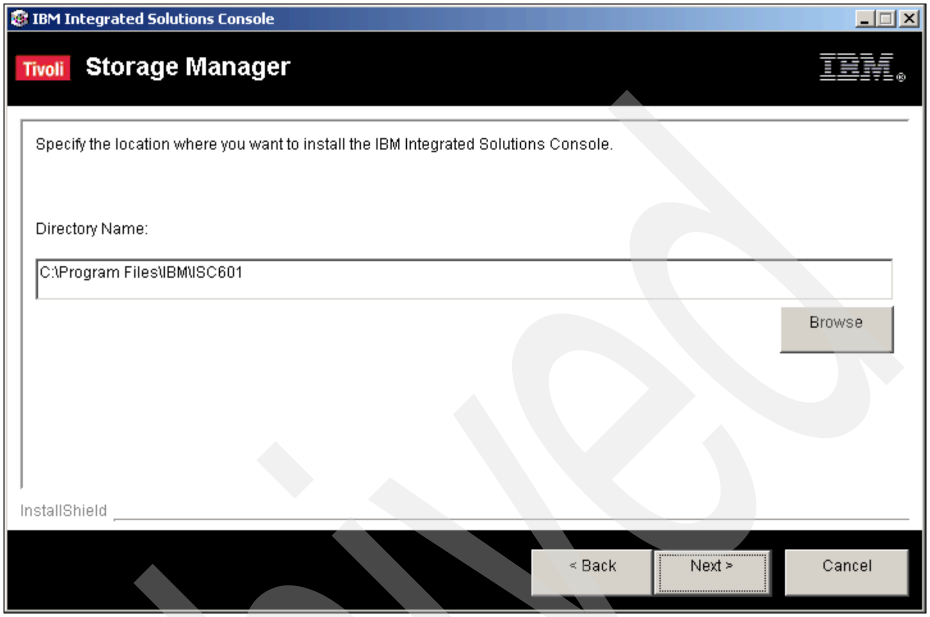

Figure 4-32 Destination selection

6. Enter port numbers for the Web administration and Secure Web administration ports, and click **Next** (see Figure 4-33).

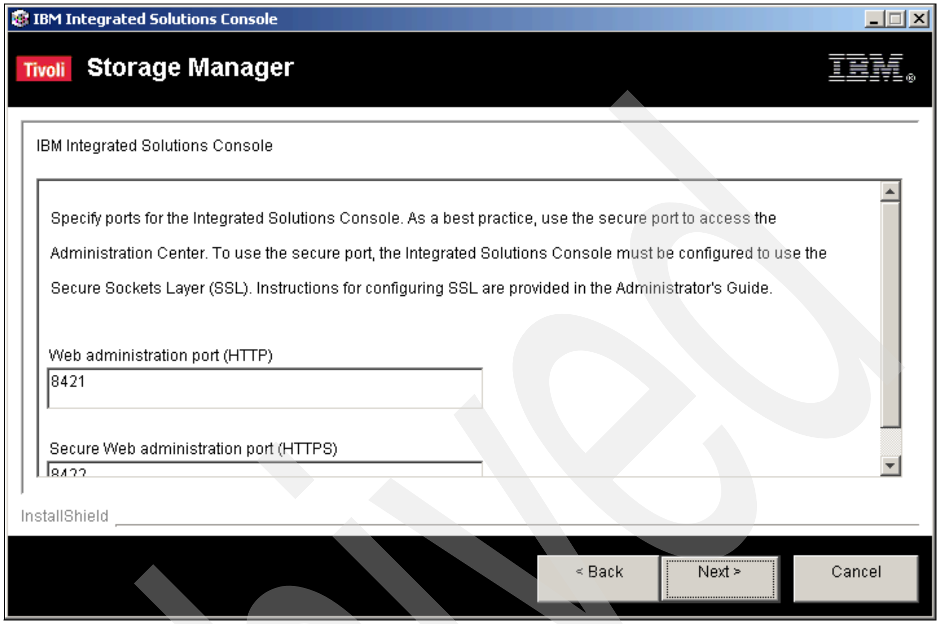

Figure 4-33 Selecting the default port for ISC

7. A review window to confirm your selections is displayed. Click **Next**.

   The installation of ISC begins. The installation takes at least one to two hours to finish, depending on the memory and processing speed of the server.

8. A completion window is displayed when the installation is complete. Click **Next**.

9. Click **Finish**

Attention: After installation completes, reboot the server before starting the Administration Center installation; otherwise the Administration Center installation fails to install.

4.6.3 Installing Administration Center

This section provides the step-by-step method of installing Administration Center on the Microsoft Windows platform. Before installing the Administration Center, make sure that you install the Integrated Solutions Console (ISC) and restart the server. The ISC must be running when you install or upgrade the Administration Center.

The Administration Center can display information in languages other than English. You can enable language support during installation. If you enable this support, the Administration Center attempts to display the same language as the Web browser used to access it. If the Web browser uses a different language than the Tivoli Storage Manager server, some information is displayed in the language used by the server. To change the default language used by the Web browser, see the browser documentation.

The installation of Administration Center is similar to the ISC install. It also begins with the Java installation.

1. Click **Next** in the InstallShield Welcome window to continue the installation (see Figure 4-34).

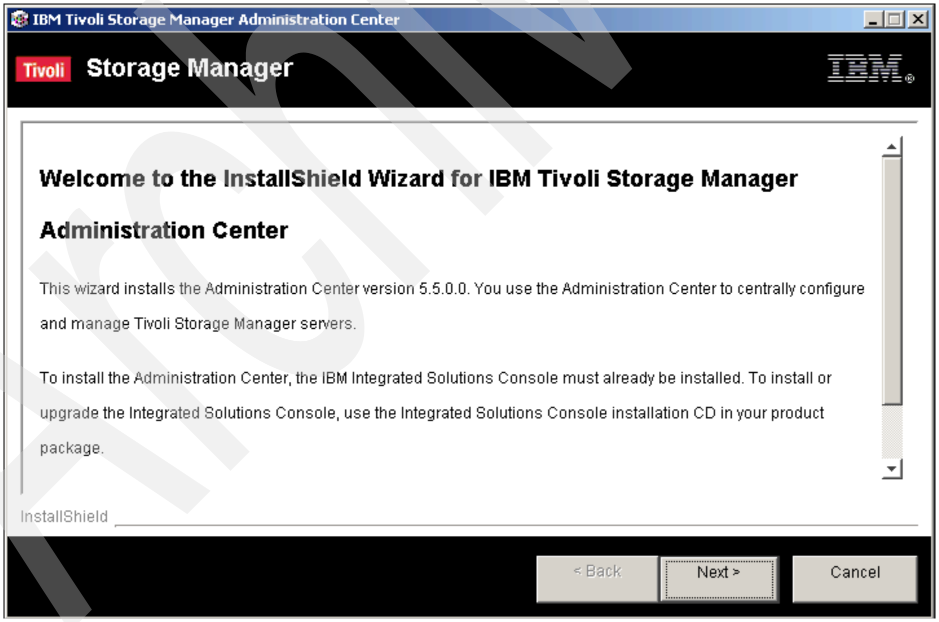

Figure 4-34 Welcome to InstallShield Wizard for Administration Center

2. Accept the International License Agreement, and click **Next** to continue.

3. The next window lists the installation path and Web administration port that you provided during the ISC installation process (see Figure 4-35). If the information is correct, click **Next**. Otherwise, click the appropriate option to update the information, then click **Next**.

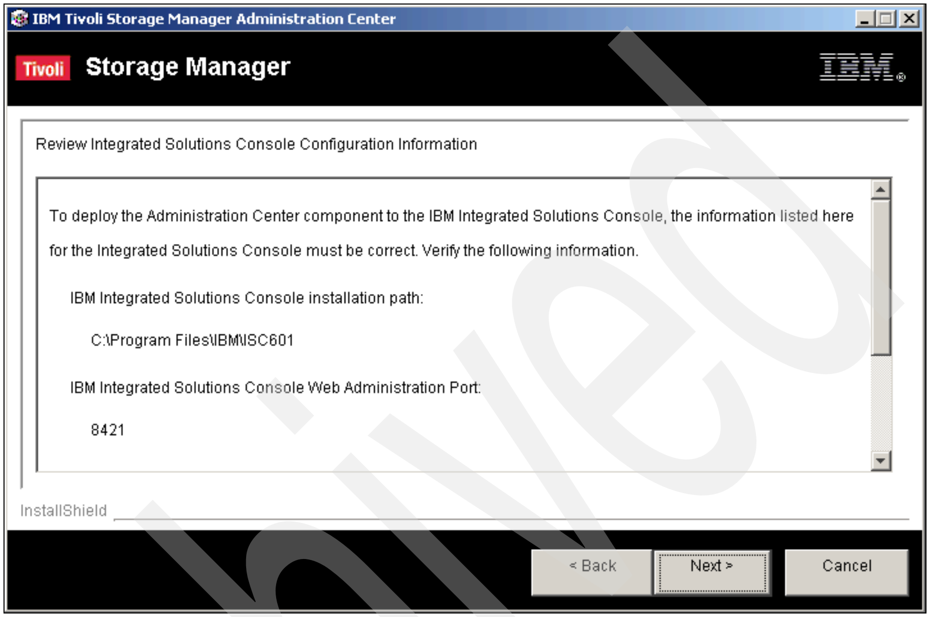

Figure 4-35 Review install path admin port and user ID

4. The ISC password is the password assigned during the ISC install. Enter and verify it, and click **Next** (see Figure 4-36).

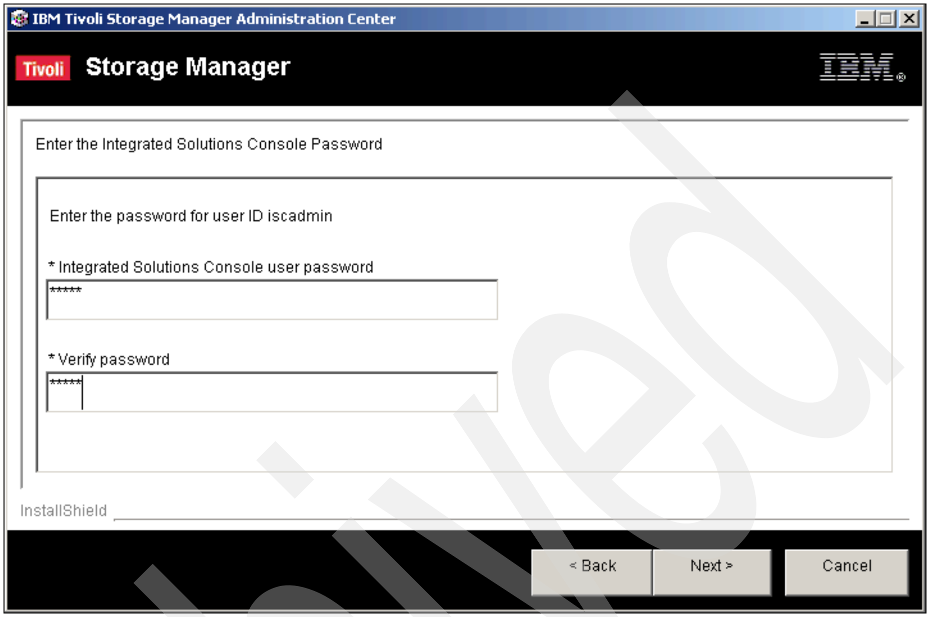

Figure 4-36 ISC password

5. If you want additional language support, select the Tivoli Storage Manager server version; if you do not require any language other than English, make no selection, and continue the installation, and click **Next**.
6. Otherwise, select the language and click **Next**.

7. For the final review of the selected installation option, press **Next** to continue (see Figure 4-37). Otherwise, you can select **Back** to make the necessary changes to the installation.

> **Important:** When the final review window is displayed, clicking **Next** starts the installation, which *cannot* be stopped when it is underway.

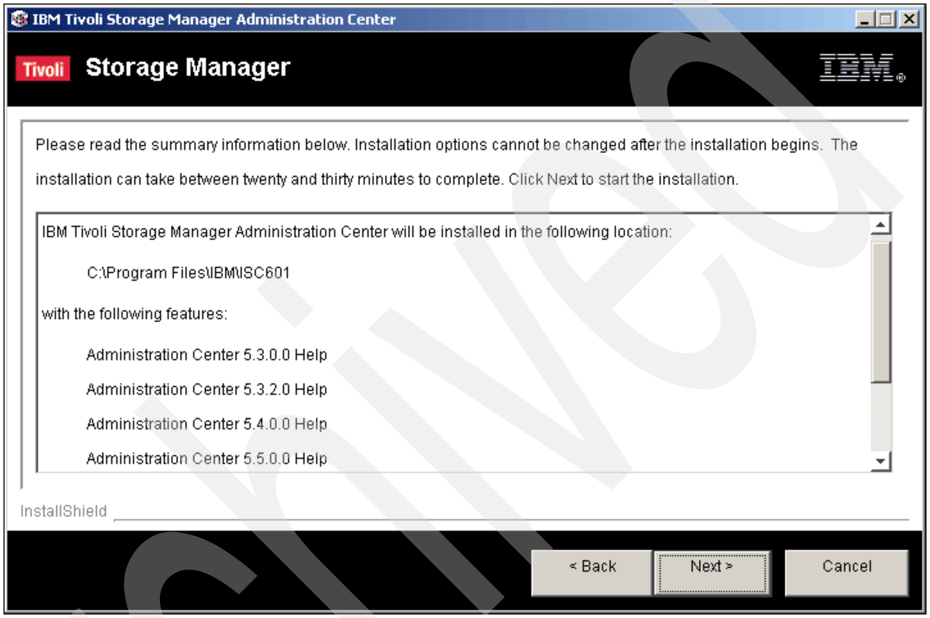

Figure 4-37 Final review of selected installation option

The next window displays the installation progress.

8. When the installation has completed successfully, click **Next** to view the installation summary (see Figure 4-38).

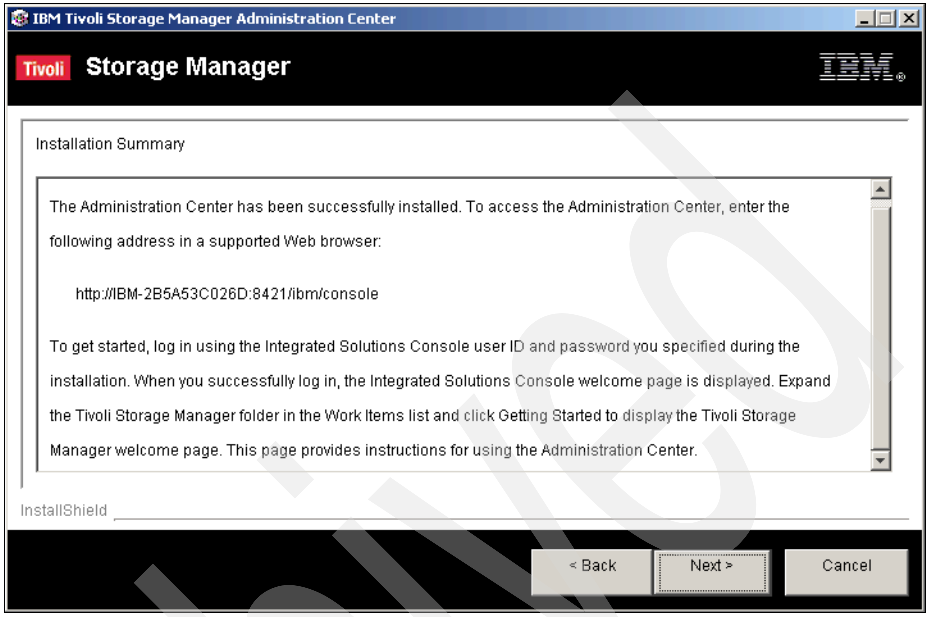

Figure 4-38 Installation summary

9. After clicking **Next**, ISC is automatically launched, as shown in Figure 4-39.

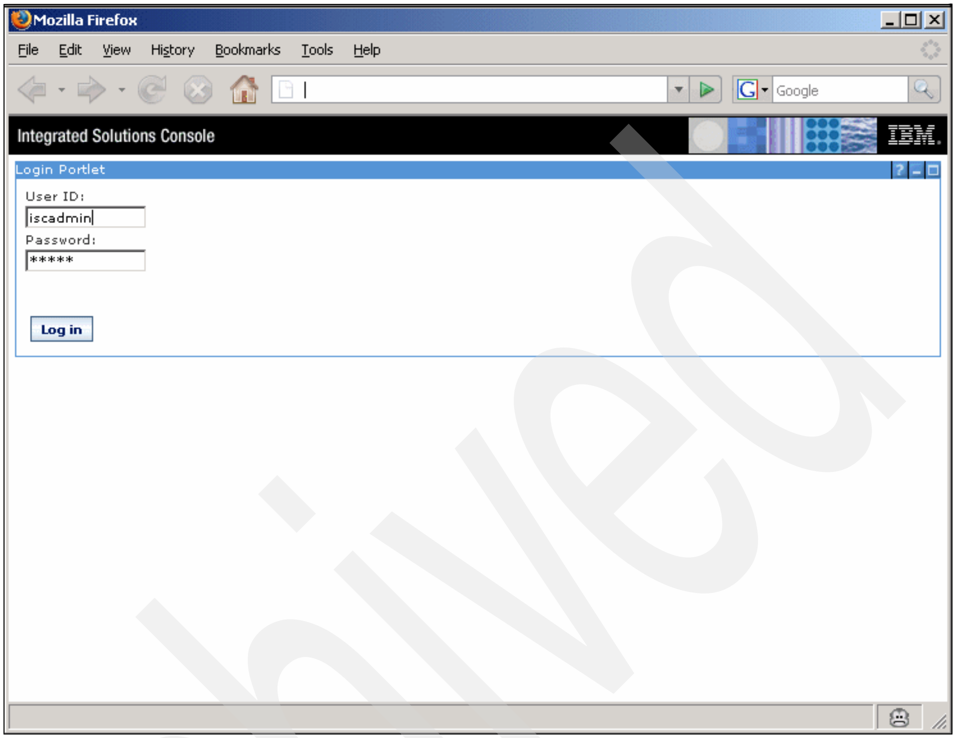

Figure 4-39 ISC logon automatically is displayed after the installation

The Administration Center installation is completed. To read more about ISC and Administration Center, refer to *IBM Tivoli Storage Manager Versions 5.4 and 5.5 Technical Guide*, SG24-7447.

4.6.4 Configuring Administration Center

After you install the ISC and Administration Center, you must add Tivoli Storage Manager servers to the portal. This section provides configuration steps for the Administration Center.

Adding a server to Administration Center

Follow these steps to add a server:

1. Open a browser and navigate to the URL `http://`*`machine_name`*`:8421/ibm/console`. Substitute *machine_name:* with the machine_name that is the network name or IP address of the machine on

which you installed the Administration Center. The default Web administration port (HTTP) is 8421. The default secure Web administration port (HTTPS) is 8422 (see Figure 4-40).

2. Enter the Console user ID and password, and click **Login**.

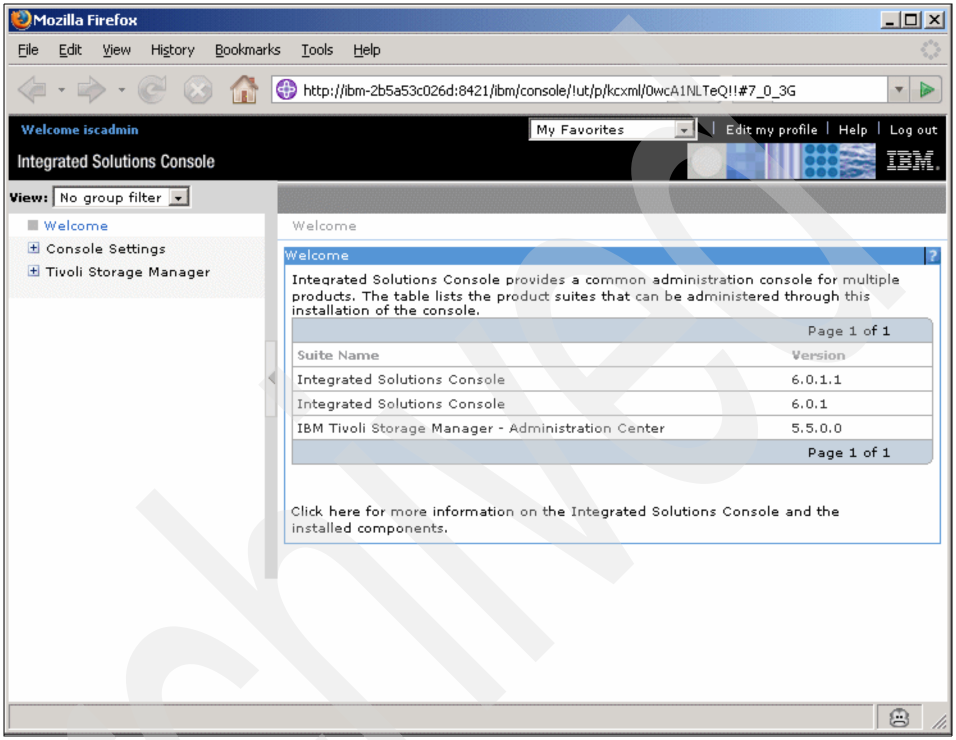

Figure 4-40 Welcome Administrative Center console

3. In the left menu tree, expand **Tivoli Storage Manager** → **Storage Devices**.
4. From the Servers window, select **Action** → **Add Server Connection** → **Go**.

5. Enter the necessary information:
   a. If this is the first installation, the Admin name and password is admin.
   b. During installation, the admin name ADMIN_CENTER is created to monitor your Tivoli Storage Manager server. To use it, you have to unlock it by selecting the check box for **Unlock the ADMIN_CENTER Administrator** (see Figure 4-41). Click **OK**.

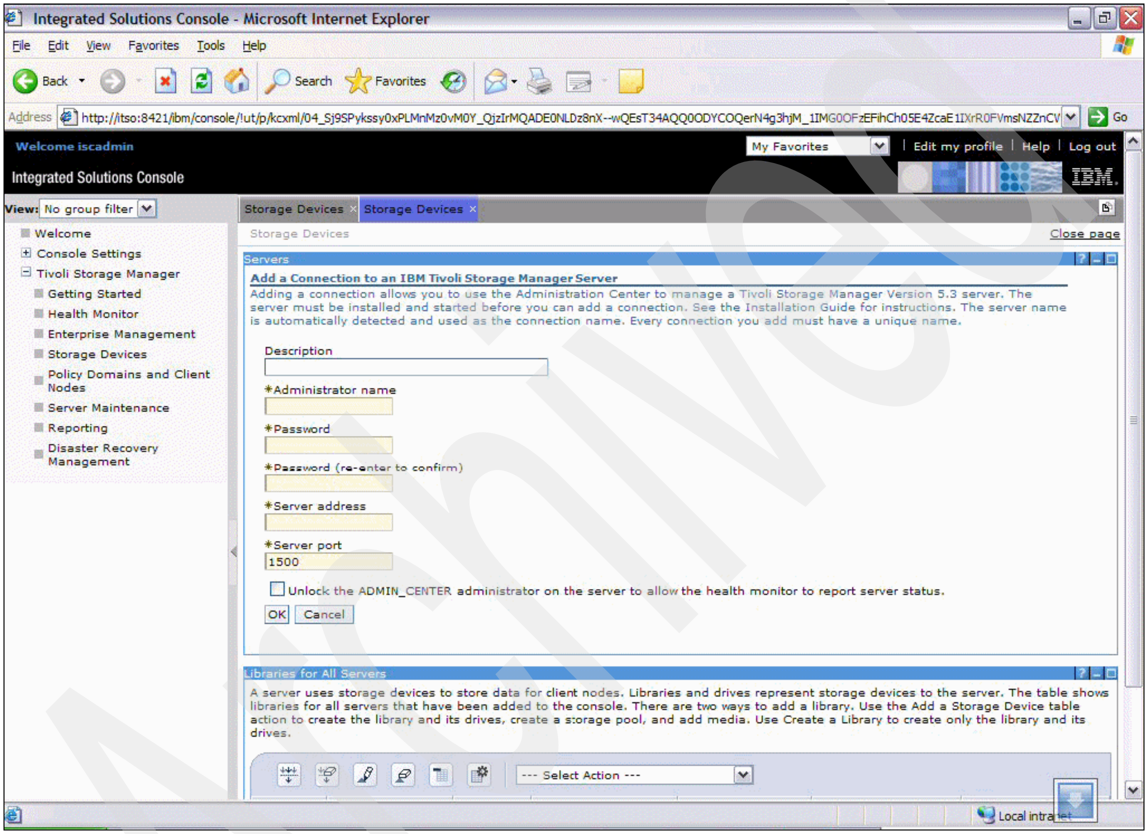

Figure 4-41 Adding IBM Tivoli Storage Manager server to Admin Center

6. The health monitor default password is ADMIN_CENTER. Change the password and, optionally, the refresh time for the health monitor (default is 10 minutes).

Your Tivoli Storage Manager server is now ready to configure.

Checking the health of the Tivoli Storage Manager server

To check the health of the Tivoli Storage Manager server, complete the following steps:

1. In the left menu tree, click **Tivoli Storage Manager** → **Health Monitor**.

   The Health column to the right of the panel displays various messages:

   – Normal (green): The server is running; the health monitor found no problems.

   – Unknown (blue): The server cannot be contacted. The following reasons may account for this status:

     - The server is not currently running.
     - Network problems prevent communications with the server.
     - The administrator name that the health monitor uses, ADMIN_CENTER, is locked or does not exist on the server.
     - The health monitor has internal errors.

   – Warning or Needs Attention (yellow): The health monitor detected conditions that might develop into significant problems. You must take preventive actions for the server's database, storage devices, or both.

   – Critical or Error (red): The health monitor detected significant problems in the server's database, storage devices, or both. The problems prevent the server from operating normally.

2. If you want to know a more detailed status for a server, click the server name for more details (see Figure 4-42).

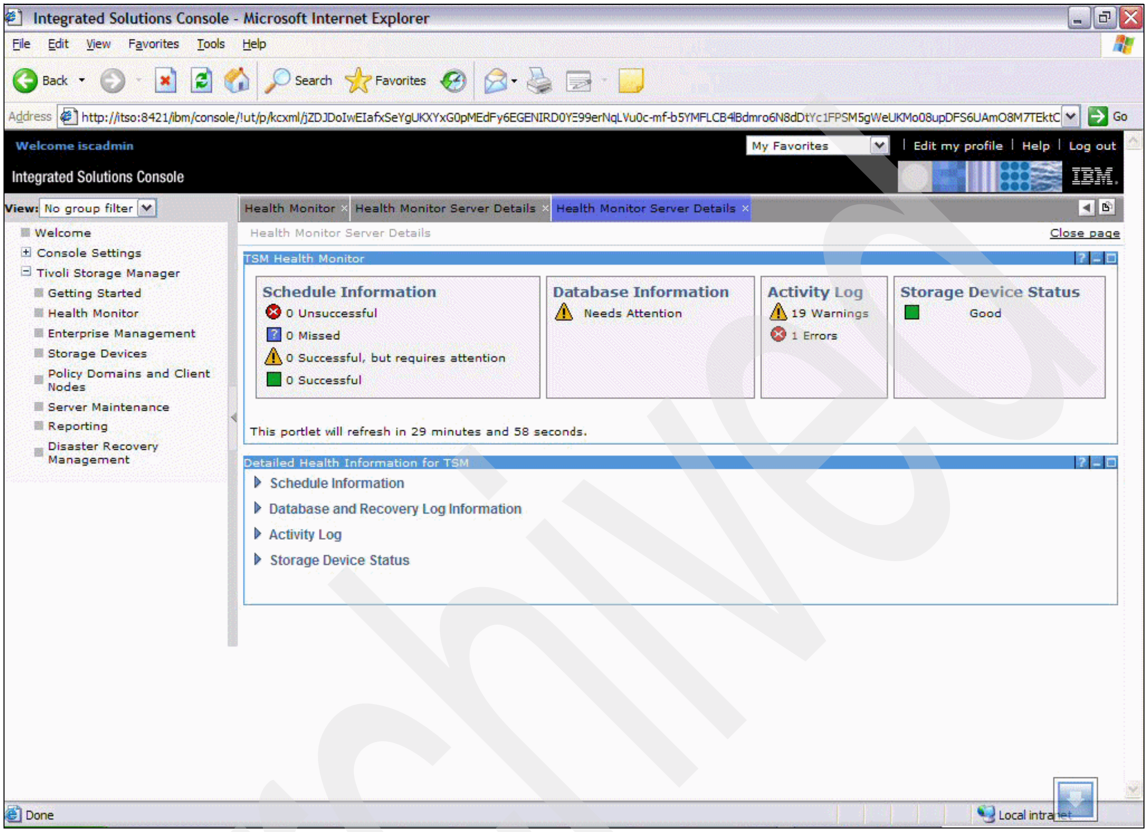

Figure 4-42 Detailed health information

The window in Figure 4-42 displays the overall status of your Tivoli Storage Manager server. Identify the problem area and select the appropriate detailed report. The Database Information section shows a Needs Attention flag. Choose the **Database and Recovery Log Information** report.

The Detailed Health report for the Database and Recovery log (see Figure 4-43) shows that it is necessary to extend the database. You do not have a backup from the last 24 hours, and the Cache Hit Ratio is less than its ideal value.

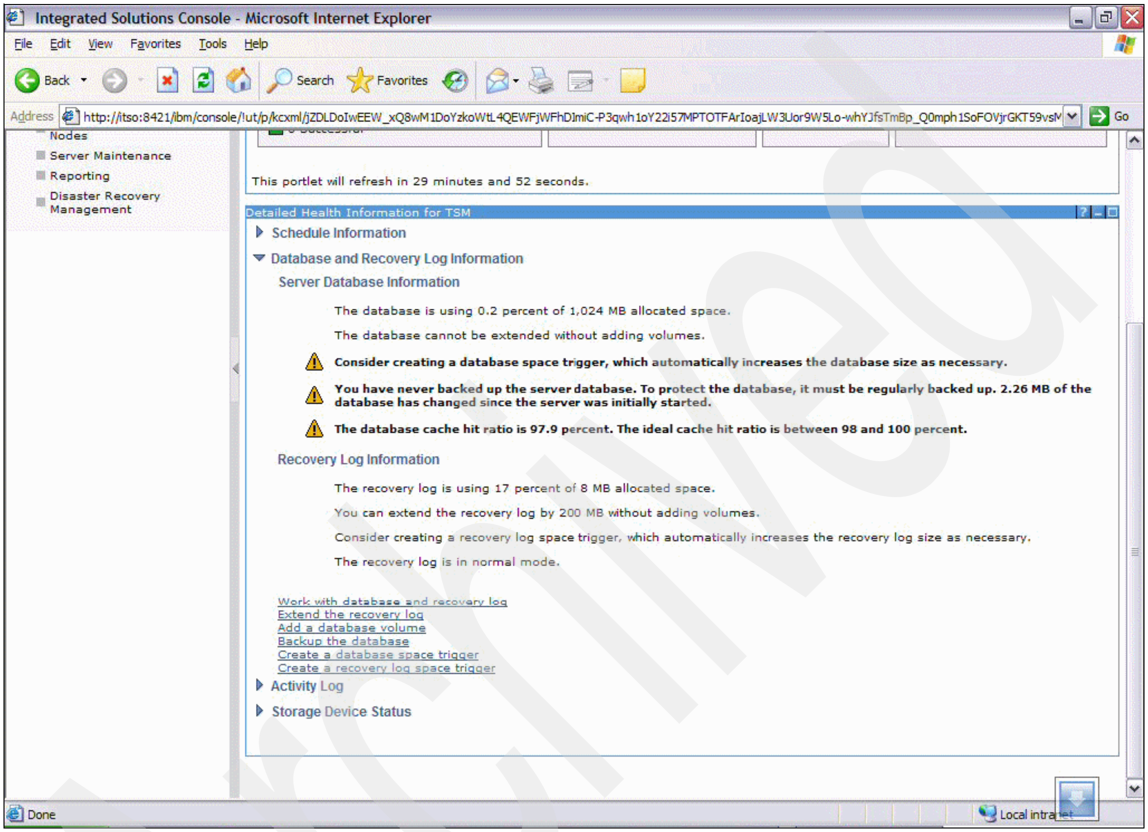

Figure 4-43 Database and Recovery log detail information

3. Perform the correct modifications, and wait for the next refresh.

Adding database and recovery volumes

Complete the following steps to add database and recovery volumes (see Figure 4-44):

1. Select **Storage Devices**.
2. Click the server name to access the Server Properties window.
3. Select **Database and Log**.
4. Select the **Add Volume** task from **Select Action** on top of the Database Section, and click **Go**.
5. Fill in the Database Volume name and capacity, and click **OK**.

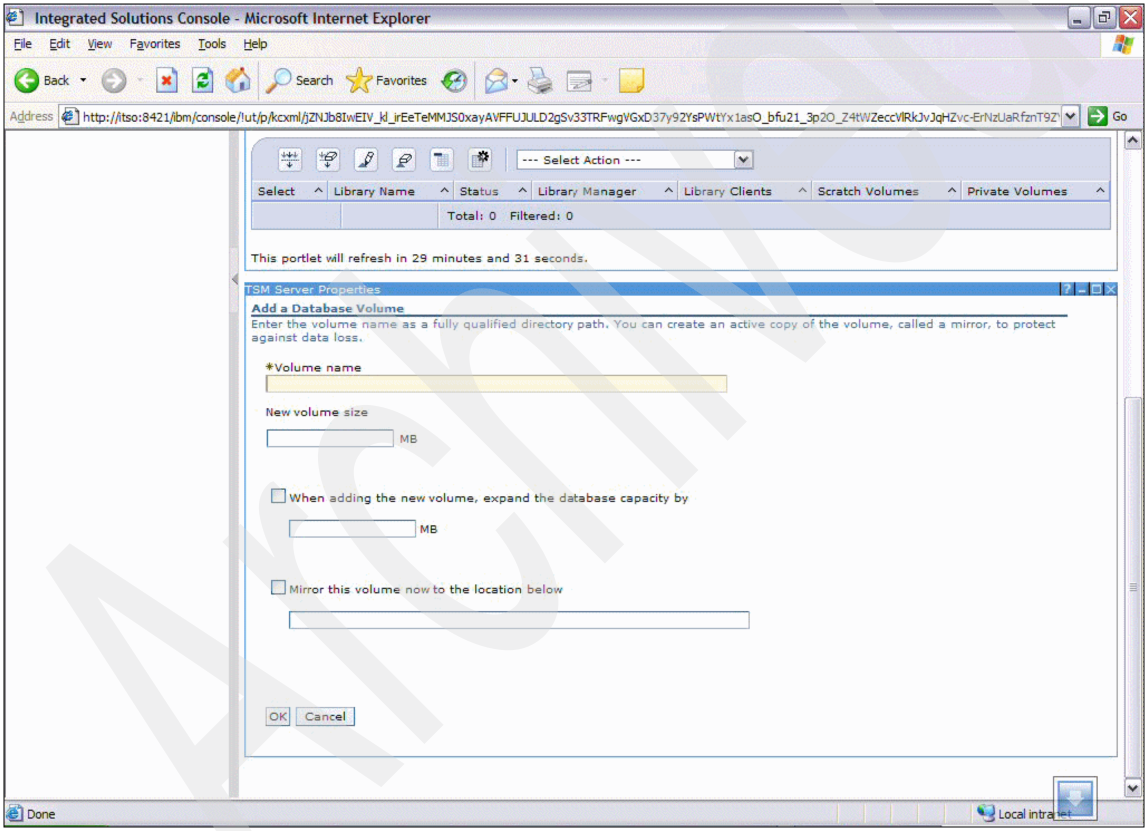

Figure 4-44 Adding database volume

With this release, it is possible to extend the database while creating it.

Note: To create a volume using something other than the system default path, enter the fully qualified directory path name in the Volume Name field.

4.6.5 Operational Reporting

In this section we briefly discuss the Operational Reporting tool. The IBM Tivoli Storage Manager Operational Reporting feature automates some of the monitoring tasks you typically perform manually. By generating reports and monitors, Operational Reporting notifies you if a server requires attention.

You can schedule Operational Reports to run daily, and the reports are generated even if no problems exist. Operational monitors are special types of reports that you can schedule to run hourly. The monitors send a notification only if issues exist. Operational Reporting does not maintain a separate database of information and is not a trending tool.

Operational Reporting is a free tool that you can use with any Tivoli Storage Manager server platform. It runs on Microsoft Windows platforms, but it can report on any Tivoli Storage Manager server.

This tool is part of the Tivoli Storage Manager for Windows server and is also available as a stand-alone package for a Windows server at no charge.

For information about installing the stand-alone package, see “Installing the IBM Tivoli Storage Manager Operational Reporting Stand-alone Package” in the *IBM Tivoli Storage Manager for Windows Administrators Guide*, SC32-0121.

You administer Operational Reporting through the Microsoft Management Console (MMC) on a Windows machine. All platforms of Tivoli Storage Manager servers, Version 5.4 and later, are supported. Operational Reporting runs as a service and supports multiple Tivoli Storage Manager servers running on a single machine.

An Operational Report consists of the following parts: a standard report, a customized summary, and optional extensions that you can create. You can select which sections to include in the report. The Operational Reporting installation package contains two default custom summary templates; one of the summary templates is for a report, and one for a monitor.

Default e-mail messages notify you if the server is running smoothly, or if such issues as failed or missed schedules are encountered. You can also link to a Web summary page to check Operational Reports about your server. An operational monitor notifies you either through e-mail or by sending an instant message to your Windows desktop. Operational Reporting can write data to a file that can be read by an IBM Tivoli Enterprise Console® log file adapter. The log file adapter reads the information and forwards it to the Tivoli Enterprise Console.

For more information about Operational Reporting, refer to *IBM Tivoli Storage Manager Versions 5.4 and 5.5 Technical Guide*, SG24-7447.

4.7 Summary

This chapter discussed the installation and configuration of the Tivoli Storage Manager and its related components. You can now perform a successful installation of the Tivoli Storage Manager server and client in a Microsoft Windows, AIX, and Linux environments.

5

Scenarios and best practices

In this chapter, we describe a typical IBM Tivoli Storage Manager deployment scenario. In this example, the firm Education Company already has implemented a backup solution based on Tivoli Storage Manager. We intend to design and expand the solution with the most important functions that IBM Tivoli Storage Manager Version 5.5 offers. Also, we describe the best practices and recommendations of each functionality. The discussion includes the following topics:

5.1 Example scenario

The Education Company provides IT courses and workshops all over the U.S., and the company also offers infrastructure classes for organizations interested in training their employees. The Education Company is a solution provider and is partners with some of the major organizations in the IT industry. Figure 5-1 shows their current IT environment.

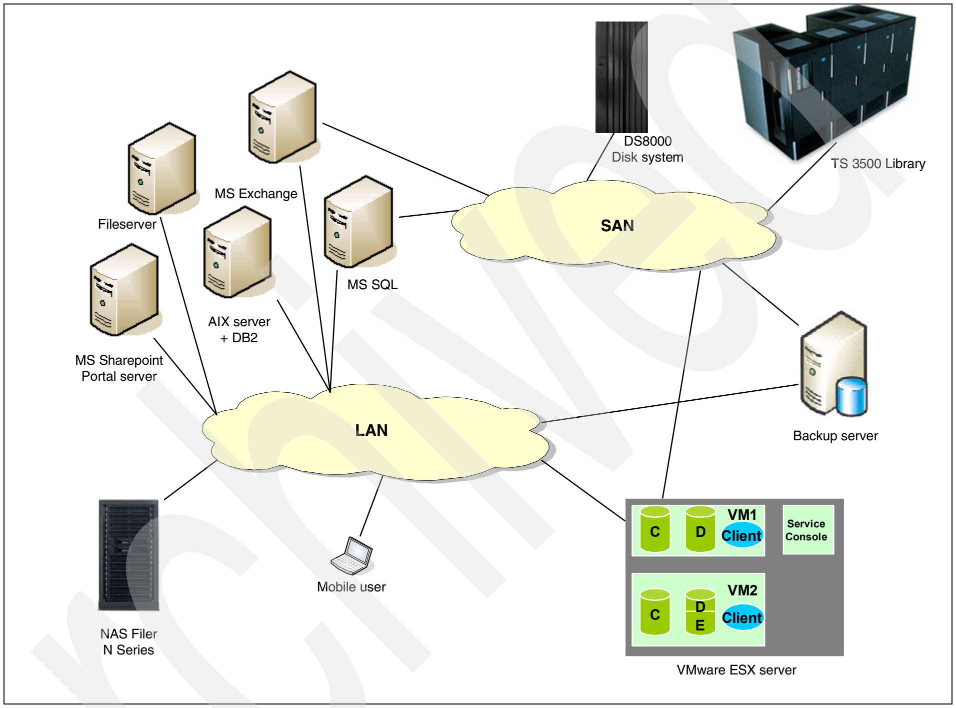

Figure 5-1 Education Company IT environment

The Education Company operates from a single data center and their environment consists of the following components:

- A Tivoli Storage Manager backup server
- VMware ESX server running virtual machines for their classes
- Application and database servers
- Network Attached Storage (NAS) filer
- Storage Area Network (SAN) systems
- Workstations and mobile users

The VMware ESX virtual machines are available over the LAN for the classroom courses and are backed up daily along with the file servers, database, and application servers during the appointed backup window. In addition, the NAS filer is backed up using CIFS or using NFS (with network-mapped drives).

The Education Company application servers are Microsoft Exchange servers and a SharePoint portal server. The file servers run Microsoft Windows and Linux operating systems and store a lot of historical data. The database servers are Microsoft SQL Server and DB2 Universal Database. The DB2 instance is installed on an AIX platform. The database server backups are taken offline.

The administration of this environment has become challenging. The Education Company is having problems with its backup window, and the overall backup performance is not optimal. The backup window frequently expires, and Education Company administrators have been manually canceling the backups. In addition, the restore is not reliable. Because many users have access to the backup server and can start backups or restores on their own, it is difficult to know who is responsible for the backup server. The library and the tapes are not managed efficiently, and the company does not have a disaster recovery plan.

5.1.1 Recommendations for implementing Tivoli Storage Manager

The Education Company wants to expand their present Tivoli Storage Manager solution and to add complementary products to address these shortcomings.

The current IT environment, from a Tivoli Storage Manager perspective, is considered a simple-to-moderate environment; the company runs only one server instance but their environment includes larger database and application servers that must be backed up.

Before installing Tivoli Storage Manager, review the following recommendations (which apply to our example and every implementation):

- Perform an environment inventory and gather information about the servers in general. In this case, the environment includes file servers, database and application servers, workstations, and mobile users that must be backed up. Also inventory the external storage and tapes, and SAN and LAN infrastructure.
- Understand the amount of data required to be backed up. You can determine the available backup window and client expectations, in terms of throughput and restores and determine your storage (disk pool size and library size) based on the total backup data.

- Calculate your throughput requirements and be sure the clients can deliver the throughput needed. You might have to study how the data is laid out on disk to improve performance.
- Planning for restore is as important as planning for backup. You do not want to realize that your system is undersized when you need a large-scale restore.
- Implement progressive-backups. Only the changed data is backed up on a daily basis. Database management backups are normally full backups. Tivoli Storage Manager completes your solution with a list of complementary Tivoli Data Protection modules, enhancing the possibility of online backups. See 1.5.5, "IBM Tivoli Storage Manager for Databases product family" on page 24 for more information.
- Determine when the backup must start and when it must end; normally the backups cannot run over into the following day.

We must plan for our Tivoli Storage Manager server instance. The core of the server is its database, storage pools, and recovery log. We recommend that you put Tivoli Storage Manager components on different devices, even different servers. All three Tivoli Storage Manager components have different I/O behavior and different cache requirements. Isolating each component on separate disk drives improves the overall performance.

For more information about how to size your database, recovery log, and storage pool, consult 2.3, "Sizing your Tivoli Storage Manager solution" on page 63. After determining the size of our database, recovery log, and disk storage pools, we can proceed to install our backup server instance. The installation instructions are described in Chapter 4, "Installation and configuration" on page 119.

5.1.2 Complementary Tivoli Storage Manager products

Finally, once you have an operational Tivoli Storage Manager server, you must address the following questions:

- Where is my data going and is it secured?
- How do I back up my database servers online?
- How do I back up the virtual machines without depleting the VMware ESX server resources?
- Does the NAS filer have a better method of backup?
- How do I use my storage area network more efficiently?

We address these questions in this chapter with our discussion of the Education Company's use of Tivoli Storage Manager and its add-on products. This chapter also explains Tivoli Storage Manager usage and new enhancements in Version 5.5.

5.2 Security plan and implementation

These days, most organizations are dependent on their IT infrastructure, in terms of automation, data processing, and storing their intellectual property. Every organization must be aware where the data is going and whether it is safe from intruders and other threats. Planning for a secure and reliable environment has become the highest concern of these organizations because their data is their most important asset.

Today's storage management has to go beyond traditional backup and recovery solutions. Data leads the e-business economy. A well-implemented storage management security infrastructure ensures data reliability, scalability, and disaster recovery, and protects the accuracy of information as a whole.

In a typical moderate-to-complex Tivoli Storage Manager environment, Tivoli Storage Manager is the principal application that reaches into every corner of the enterprise, from the largest database server to the desktop. Not only is the Tivoli Storage Manager server likely to have the most recent backups from many systems, but it likely has large quantities of historical data, potentially going back for years. The Tivoli Storage Manager server has the ability to execute commands and control applications on the clients attached to it.

Tivoli Storage Manager must therefore be viewed as one of the most valuable and powerful applications in any organization where it is widely deployed. As it is, Tivoli Storage Manager must be secure.

In this section we describe ways to protect your data using Tivoli Storage Manager.

> **Note:** Tivoli Storage Manager security is extensively covered in *IBM Tivoli Storage Manager: Building a Secure Environment*, SG24-7505.

5.2.1 Introduction to encryption

Encryption is the translation of data into a secret code and is the most effective way to achieve data security. To read an encrypted file you must have access to a secret key, or password, that enables you to decrypt it.

Encryption can be implemented in hardware or software:

- Hardware-based encryption

  Hardware-based encryption is performed by special processors in certain types of hardware (such as a tape drive, a network router, or a specialized appliance) that are very fast and can encrypt in real time. The data transfer rate is not impacted, and no CPU resources on the computer server are required because the encryption is off loaded to other hardware.

- Software-based encryption

  Software-based encryption uses the server CPU to do the work. Because most computer CPUs are not specialized for encryption, software-based encryption can take longer and consume CPU resources that otherwise are available for other operations.

Symmetric and asymmetric are the two main types of encryption:

- In symmetric encryption, the same key is used to encrypt and decrypt the message.
- In asymmetric encryption (also called *public-key encryption*), one key encrypts a message, and another key decrypts the message.

A public-key cryptographic system uses a pair of unique keys (a public key and a private key). Each user is assigned a pair of these keys to encrypt and decrypt information. A message encrypted by one of these keys can be decrypted only by the other key in the pair. The keys have the following additional characteristics:

- The public key is available to others for use when encrypting information that is sent to another user. For example, people use a person's public key to encrypt information they want to send to that person. Similarly, people can use the user's public key to decrypt information sent by that user.
- The private key is accessible only to an individual. The individual can use the private key to decrypt any messages encrypted with the public key. Similarly, this user can use the private key to encrypt messages, so that the messages can be decrypted only with the corresponding public key.

Tivoli Storage Manager supports both software encryption and hardware encryption. Software encryption is implemented within the Tivoli Storage Manager client. Hardware encryption is available with certain tape drives.

5.2.2 Current security environment

The Education Company is facing security issues. The major issue is the backup window. Many regular users have access to the backup server and can start backups and restores at any time. The backups tend to overlap, and no plan exists for the use of tapes or disk; the Education Company has not considered transferring large database backups to tape and file server backups to disk for faster restores and availability.

Manageability of the tapes is inefficient because it is difficult to determine which tapes are effectively in use or inside the tape library. In addition, library capacity is inadequate. Users are currently checking out tape volumes to make room in the library to check in empty volumes for the nightly backups. This manual control, inventory, and transportation of the tape volumes results in tapes being damaged, lost, or insecurely accessed.

The Education Company provides training every day, and students who are granted access to the virtual machines can acquire confidential data. Network traffic is also a concern. The corporate network is used for training purposes, which is a security issue that we must address.

The backup server is not part of a firewall environment, and communication between the server and its clients is not protected from packet sniffers. As noted earlier, the Education Company selected Tivoli Storage Manager as its backup server.

The Education Company must implement a secure environment, which meets the following requirements:

- Securing client/server communication
- Carefully administering user access privileges
- Providing secure data transmission
- Securely implementing the Tivoli Storage Manager server
- Protecting data on tapes during transportation

We discuss these security requirements in the sections that follow. See also *IBM Tivoli Storage Manager Versions 5.4 and 5.5 Technical Guide*, SG24-7447, for a more detailed explanation of these features.

5.2.3 Client/server communication

The Tivoli Storage Manager server and client use TCP/IP sessions to communicate with each other. Tivoli Storage Manager uses a “mutually suspicious” algorithm to establish authentication between the client and server in a way that the password credentials are encrypted in a simple message, known as a *buffer*. To ensure secure client/server communications, the password is never part of the exchange between the client and the server.

Basically client/server communication involves the following two types of sessions:

- Client sessions

  These sessions, which are the typical backup/archive sessions, are between the Tivoli Storage Manager server and client. You use these sessions for transmitting data, along with the associated metadata from or to the client, to be backed up, restored, archived, and retrieved. A client usually generates two sessions; one session queries the server, and one sends file data for a backup or archive session.

- Non-client sessions

  You use these sessions, for example, when acquiring a drive from the library manager, running administrative commands using **dsmadmc**, sending SNMP traps, communicating between the Tivoli Storage Manager server and storage agents, notifying subscribers, or sending events to the event server.

Tivoli Storage Manager exploits the multithreading capabilities of modern operating systems by transparently initiating multiple backup/archive or restore and retrieve sessions on the client for rapid processing and data transfers between the client and the server.

The underlying multithreading model used by Tivoli Storage Manager is called the *producer-consumer* or *reader-writer* model. This model usually involves two basic types of threads:

- Producer (or reader) thread that writes data to a buffer, queries the server, and scan the local file systems
- Consumer (or writer) thread that reads data from a buffer and performs backup and archive

The multisession function involves five types of threads: main, signal waiting, producer, consumer, and performance monitor.

All data sent to Tivoli Storage Manager storage during a backup or archive session occurs within the bounds of a transaction. Files are not sent to the server as individual objects; instead, Tivoli Storage Manager combines multiple files in one transaction to reduce overhead and to increase performance.

When the client starts sending or receiving data, it pays attention to both sides of the communication (the server and the client). All operations are controlled in such a way that Tivoli Storage Manager can detect any data inconsistency during the transfer (due to a network problem, full hard drive, or a file that already exists, for example). This control provides a high level of data integrity for Tivoli Storage Manager. A single transaction is an atomic action, the smallest possible unit of work. Data sent within the bounds of a transaction is either committed completely to the system at the end of the transaction or rolled back if the transaction is ended prematurely.

For more detailed information regarding client sessions, refer to *IBM Tivoli Storage Manager: Building a Secure Environment*, SG24-7505.

Client communication processing

The Tivoli Storage Manager client has a variety of ways to connect to the Tivoli Storage Manager server. Each method communicates to the server by opening a session to authenticate itself. After installing the Tivoli Storage Manager client, you must configure the client before performing any operations. When you configure your Tivoli Storage Manager client, the first step is to create the client options file (dsm.opt for Microsoft Windows and dsm.sys for UNIX and Linux). In this client options file, you identify the server to connect to and the communication protocol to establish the communication. The client options file can also include authorization, backup and archive processing, scheduling, and other options that dictate the client's behavior.

The communication method is one of the most important purposes of the client options file. You choose the communication method and the port in the COMMMETHOD parameter and the TCPPORT parameter, respectively; the communication method is primarily TCP/IP, and the port ranges from 1000 to 32767, with a default of 1500. You specify the server to connect to in the TCPServeraddress parameter.

When a native backup/archive client logs onto the server using its node's name and password, or an administrator's ID and password, the user ID that is provided by the user is checked against the node name. If they match, Tivoli Storage Manager attempts to authenticate the session as a node name. If they do not match or the authentication fails, the server attempts to find a registered administrator to compare with the ID provided by the user and authenticates the session with this ID and password.

Tivoli Storage Manager is intended to be used in a relatively secure environment. It is assumed that an organization must have external firewalls and physical security in place at a minimum.

Figure 5-2 depicts the data flow between the client and the server.

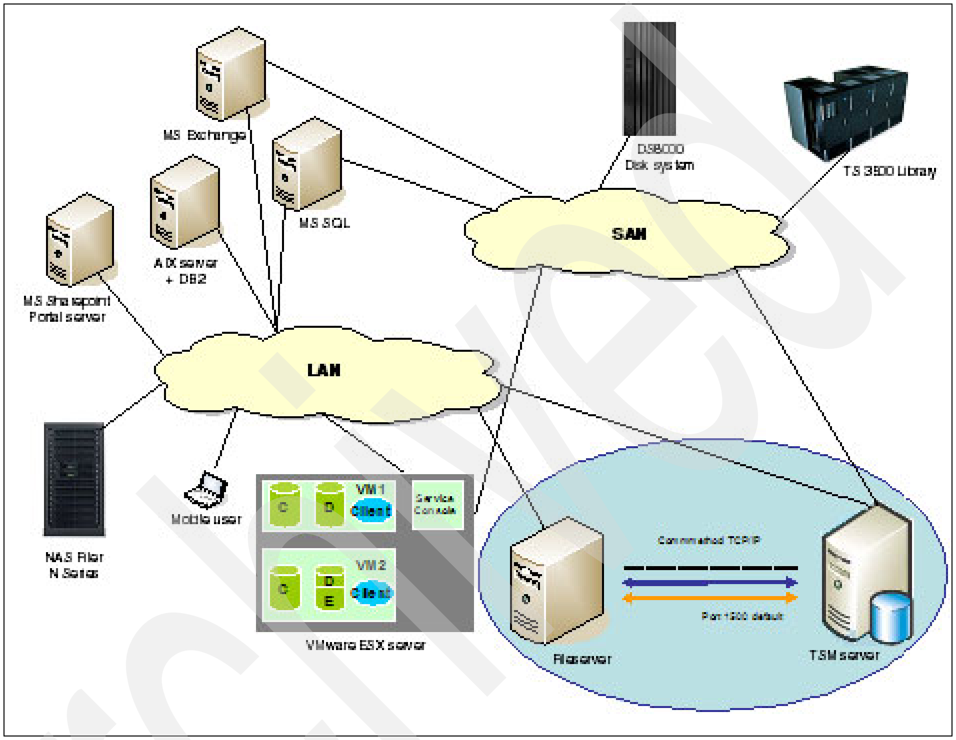

Figure 5-2 Tivoli Storage Manager client/server communication

Follow these recommendations:

- Change the port number from the default port 1500, especially if you have your Tivoli Storage Manager server in the DMZ (demilitarized zone).
- Restrict client sessions to a specific port number and configure your server to not accept client-initiated sessions.
- The backup/archive client communicates with the server using the command-line interface `dsmc`, Java GUI interface, Web interface, or through the Scheduler service. The methods for implementing security in these utilities are covered in Chapter 4, "Securing the client," in *IBM Tivoli Storage Manager: Building a Secure Environment*, SG24-7505.

Password handling

The two methods of accessing passwords are through the PASSWORDACCESS option in the client options file, which determines which method to use, prompt or generate. Use the following two methods to access passwords:

- If you set the PASSWORDACCESS option to prompt, the user is prompted to enter the user ID and password interactively every time. This method does not allow unattended operations because the Scheduler service usually runs in the background.
- If you set the PASSWORDACCESS option to generate, the user establishes the first sessions with the server using either **`dsmc`** or **`dsm`**, and the password is encrypted and stored locally in the Tivoli Storage Manager.PWD file on UNIX, Linux, NetWare, and Mac systems. On Microsoft Windows systems, the password is stored in a registry key. When the password expires, a new password is generated automatically and encrypted, and it is stored again, replacing the previous password.

It is not unusual for a single client system in a configuration to have multiple nodes to access the server (using different node names). The multiple nodes differentiate the retention period for each daily, weekly, monthly, and annual backup that is taken and can be associated with many Tivoli Storage Manager servers.

An affinity is kept on the password file that associates the encrypted password with the node to which it belongs. This affinity is slightly different on Microsoft Windows, where the passwords are kept in the registry, rather than in a physical file as it is in UNIX and Linux systems.

You can apply password rules. Some parameters in the Tivoli Storage Manager server help in hardening passwords and access policies. For more information see Chapter 4, “Securing the client,” *IBM Tivoli Storage Manager: Building a Secure Environment*, SG24-7505.

Access controls

Many factors are involved in determining the resulting access privileges that a user is granted when dealing with Tivoli Storage Manager.

Two types of user or authorization processes take place on the client system. The first process relates the operating system rights, as well as file and directory ownership and permissions. The second type of authorization process happens in the Tivoli Storage Manager realm.

Each backed up, or archived, object has an owner recorded in Tivoli Storage Manager, where possible authorized nodes or administrators are also recorded.

We define the following three types of users when dealing with local system rights:

- Superuser

  The administrator user on UNIX systems is also known as the root user and has an ID with the value of 0 (zero). A user with the user ID of 0 can perform any operating system task and manage Tivoli Storage Manager client resources. Microsoft Windows systems have the same concept, which uses the user name *Administrator* and a local group called *Administrators*. After a user is included as a member of this group, the user has all privileges on Tivoli Storage Manager client resources.

- Authorized user

  A user who is not a superuser but has full control of Tivoli Storage Manager client files is referred as an *authorized user*. For UNIX systems, this user is usually the owner of the files that make up the client package.

- Unauthorized user

  A *unauthorized user* means a user who is neither a superuser nor the owner of the Tivoli Storage Manager client files and folders but has permission to access Tivoli Storage Manager client services.

You have to determine which type of user in your environment is able to run backups and archives and then determine the type of access they must have to the operating system. The Tivoli Storage Manager access level is dependent on the access users gain when they log onto the operating system.

In the Tivoli Storage Manager server, the server's administrator can control the role assigned to the administrator. As a Tivoli Storage Manager administrator, you can use the command GRANT AUTHORITY to grant privileges to an administrator and the command REVOKE AUTHORITY to remove those privileges.

For UNIX and Linux clients, you can restrict or allow user access to the Tivoli Storage Manager client package, which provides higher flexibility to use and manage the system. However, this flexibility increases the risk of undesirable data exposure resulting from a mistake in the configuration tasks.

The most restrictive environment allows Tivoli Storage Manager access only from the root account. Other users are denied access to the executable binaries and configuration.

For Microsoft Windows clients, regular users can create a unique options file and include the options SCHEDLOGNAME and ERRORLOGNAME, which must point to a directory where the user has write access. Users can then call the **dsm** or **dsmc** executable using the options file they created.

However, to be able to perform backup and restore operations of an entire machine, the user must be a member of either the Administrators group or the Backup Operators group. With the privileges granted by one of these groups and assigned rights to the Manage auditing and security log, the user can access all files across the system. Using these privileges, you can restrict or allow users to perform backups and restores.

When registering a node, you do not have to provide an administrator for that node. By default, an administrator with the same name as the node name is created. This administrator is automatically granted the node's owner privilege for that node.

You can decide whether to use a separate administrator ID for each client node at the time of registration. Regular users require this ID for remote access through the Web client. From a security perspective, the administrator must open only the necessary ports. If you do not have to give a user this type of access for a specific node, you can use the parameter USERID=NONE or explicitly define a separate administrator for this node at time of registration.

If regular users are allowed to perform backup, archive, restore, and retrieve operations on their own using the Web client, the best approach is to create administrator IDs using their user names or other identification that is linked to the person, not the machine. These user administrators are then granted the client owner authority to the nodes that they are allowed to operate.

If this access is not required, as the Tivoli Storage Manager administrator, you can prevent this access by not creating this administrator ID and by allowing other Tivoli Storage Manager administrators to perform these operations at the request of the regular users

For more information, see Chapter 4, "Securing the client," *IBM Tivoli Storage Manager: Building a Secure Environment*, SG24-7505.

5.2.4 Client data encryption

Data security is a major concern for most organizations. Tivoli Storage Manager Version 5.5 addresses security-related requirements. When Tivoli Storage Manager client encryption is used, data is encrypted by the client itself, and data is therefore protected as soon as it leaves the client for its entire life span in the Tivoli Storage Manager storage hierarchy.

Most Tivoli Storage Manager clients, at V5.4.0 or higher, can encrypt data using either DES56 or AES128 encryption. Information regarding these encryption standards is available at the National Institute of Standards and Technology Web site:

http://csrc.nist.gov/CryptoToolkit/tkencryption.html

Both DES and AES are block ciphers and use symmetric key algorithms, which require substantially less CPU power than asymmetric algorithms. Symmetric key encryption schemes require the use of a shared secret key.

Prior to Tivoli Storage Manager V5.5, encryption using the backup/archive client required either that the user remember the encryption key password during restore or that the password be stored locally on the client system. Tivoli Storage Manager V5.5 is enhanced so that Tivoli Storage Manager generates, encrypts, stores, and manages the encryption key in the Tivoli Storage Manager database. See 5.2.5, “Client transparent encryption” on page 222 for more information.

The Tivoli Storage Manager backup/archive client can be used to encrypt data, with the client user performing key management. Files to be encrypted are selected using the include.encrypt and exclude.encrypt processing options. The backup/archive client supports two options, ENCRYPTKEY=PROMPT (default) or ENCRYPTKEY=SAVE.

> **Note:** Consult the *Backup-Archive Clients Installation and User's Guide* for your platform for more information about the include and exclude options. You can access these guides at:
>
> http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp

When using ENCRYPTKEY=PROMPT, Tivoli Storage Manager prompts you to specify the key upon every invocation of the backup/archive client. You therefore have to remember each client key you use because the key is not saved in the Windows registry or the Tivoli Storage Manager password file on UNIX. If you forget the key, your data is unrecoverable.

When you use rn, your specify that the encryption key is saved to the local password file (TSM.PWD in UNIX, Linux, Macintosh, and NetWare, and on the Windows registry) for use on subsequent invocations. All files backed up or archived by that node then have the same encryption key.

Based on the method of encryption for the backup/archive client, we describe the following backup and restore processes:

- Backup process:

  After authenticating, the user starts a **dsmc** session. The Tivoli Storage Manager client reads the dsm.opt and using the include.encrypt option, the client encrypts all data before sending to the server. The client checks for the parameter ENCRYPTKEY. If it is set to prompt, the user is prompted for the encryption password. The user must remember the encryption password in order to restore data. If it is set to save, the password is then stored locally in the Microsoft Windows registry, or in the TSM.PWD file for UNIX and Linux, and is not requested for further backups or restores.

- Restore process:

  After authenticating, a user starts a **dsmc** session to request data to the Tivoli Storage Manager server. If the ENCRYPTKEY parameter is set to prompt, the user is prompted for the password before proceeding. If it is set to save, the client uses the password either in the Windows registry or in the TSM.PWD file for UNIX or Linux.

The data is unencrypted by default, but the sessions between the client and the server are encrypted by default. Encryption can be easily configured on the client side with the PASSWORDACCESS and ENCRYPTKEY parameters described previously.

Figure 5-3 shows the backup/archive client encryption method.

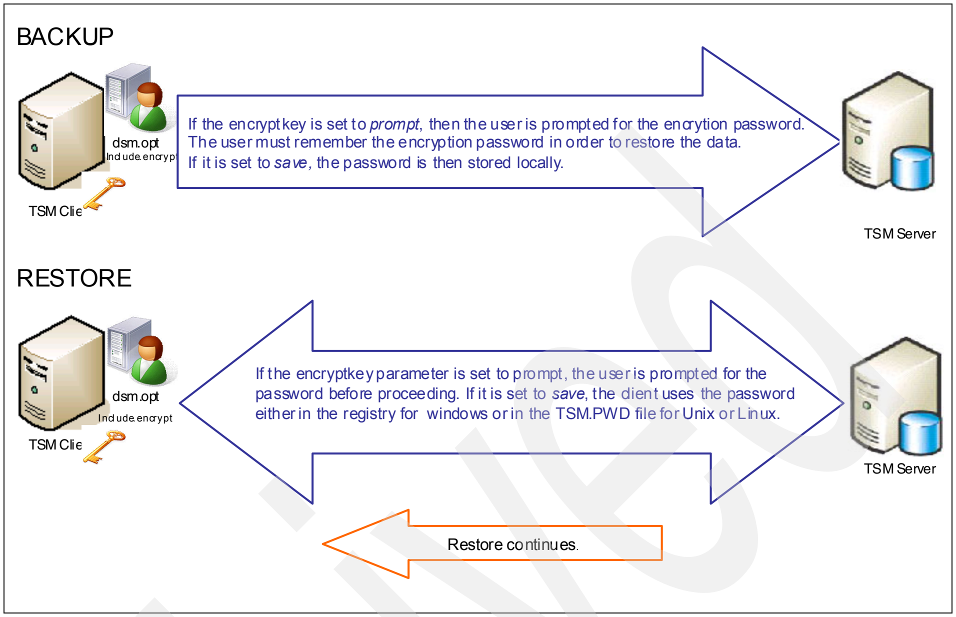

Figure 5-3 Tivoli Storage Manager backup/archive client encryption method

Using client-side encryption on the Tivoli Storage Manager client can affect performance due to the process of encrypting data. This overhead is typically small, and using client-side encryption offers benefits.

The most important advantage of enabling client-side encryption is that it ensures protection from packet sniffers because data that is sent to the server is encrypted and therefore protected during transmission. This results in the encrypted data being stored in the Tivoli Storage Manager server storage pools, where it is secure and safe from unauthorized access.

5.2.5 Client transparent encryption

Tivoli Storage Manager Version 5.5, acting as a key manager, transparently generates, encrypts, stores, and manages the encryption key for the Tivoli Storage Manager backup/archive and API client backups and restores.

API data encryption was introduced in Tivoli Storage Manager client Version 5.3. All Tivoli Storage Manager complementary products capable of sending data to or from the Tivoli Storage Manager server using the API can use client encryption. The IBM Tivoli Storage Manager HSM for Windows client also uses API data encryption.

The following two client encryption methods are available:

- Application-managed:

  The application that uses the API sends the key password to the API, and then the application manages the key password. The application provides the key password in the **dsmInitEx** call and must provide the correct password at restore time to retrieve the data. The API itself does not issue a prompt to the user but relies on the application to prompt the user as necessary.

- Server-managed:

  The server-managed (or transparent) encryption method does not require any changes to the API application code in order to function. The application has no awareness that the data backed up is encrypted; the API and the Tivoli Storage Manager server handle all aspects of the encryption. A random encryption key for each session is generated, and the keys are stored in the Tivoli Storage Manager server database. Keys are encrypted in transit between the client and the server.

The API supports the ENABLECLIENTENCRYPTKEY option with a default of no. With this option set, the API internally generates an encryption key and saves the key to the Tivoli Storage Manager server. The ENABLECLIENTENCRYPTKEY option is not applicable to the backup/archive client. The option is only for an API client and is ignored by the client during backup or archive.

In Tivoli Storage Manager Version 5.5, the ENCRYPTKEY=GENERATE option applies to both the backup/archive client and the API.

The ENCRYPTKEY=GENERATE option enables transparent (server-managed) encryption. The ENCRYPTKEY option is the preferred way to specify the type of encryption to be used.

Based on the new method of transparent encryption, the following processes for both backup and restore for the backup/archive client are the same for the API:

- Backup process:

  After authenticating, the user starts a **dsmc** session. The Tivoli Storage Manager client reads the dsm.opt and using the include.encrypt option, the client encrypts all data before sending it to the server. It checks for the parameter ENCRYPTKEY. If it is set to generate, an encrypted password is automatically generated and stored in the Tivoli Storage Manager database.

- Restore process:

  After authenticating, the user starts a **dsmc** session to request data to the Tivoli Storage Manager server. If the ENCRYPTKEY parameter is set to generate, the encrypted password is returned to the client to enable the user to transparently decrypt the file for restore and retrieve operations.

Figure 5-4 shows the transparent encryption method.

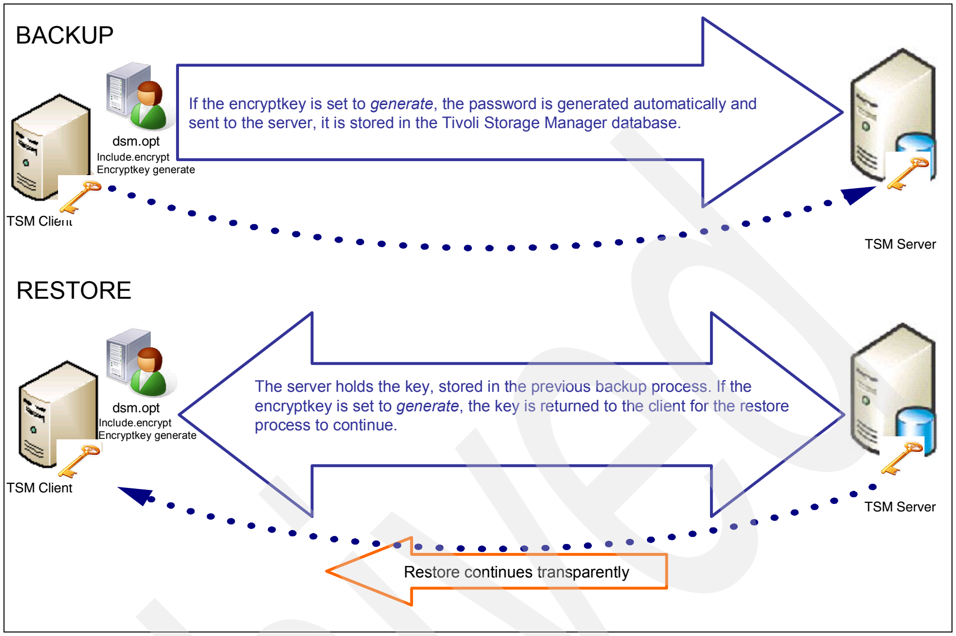

Figure 5-4 Transparent encryption method

The ENCRYPTKEY=GENERATE parameter is supported only in Version 5.5 or higher of the Tivoli Storage Manager backup/archive client.

After backing up an object with transparent encryption enabled, you can no longer restore that object with a previous version of the client. Previous clients are not capable of dealing with the object's encryption method and key. The same applies to previous server versions; objects are skipped during the process.

The encryption process affects the server processes. For more information, refer to Chapter 21, "Version 5.5 client supported environments," *IBM Tivoli Storage Manager Versions 5.4 and 5.5 Technical Guide*, SG24-7447.

The most important advantage of the ENCRYPTKEY=GENERATE parameter support in Version 5.5 is ensuring that encrypted data is sent to the server and therefore protected during transmission from packet sniffers. This process results in the encrypted data being securely stored in the Tivoli Storage Manager server storage pools, preventing unauthorized access. The other advantage of this support is that the encryption key is saved on the server. If the original machine is rebuilt, or you are trying to recover on another machine, the data is transparently unencrypted and restored on the target machine because the key is part of the server data. You do not have to access an encryption key on the original machine nor is the user running the backup required to remember it.

For more information about client transparent encryption, refer to Chapter 5, "Client data encryption," *IBM Tivoli Storage Manager: Building a Secure Environment*, SG24-7505.

5.2.6 Using Secure Sockets Layer communication

Tivoli Storage Manager Version 5.5 introduces another level of data protection by using Secure Sockets Layer (SSL) for communication between the server and the backup/archive client, API, and administrative command-line client (**dsmadmc**). SSL, which was developed by Netscape, is the standard technology for creating encrypted links between servers and clients. SSL provides a secure channel for servers and clients to communicate over open communications paths. With SSL, the identities of the parties are verified using digital certificates.

The IBM Global Security Kit (GSKit) is a required component for SSL. Tivoli Storage Manager server and client installation procedures automatically and silently install the GSKit. SSL communication encryption is available for AIX and Microsoft Windows platforms.

Figure 5-5 depicts SSL configuration.

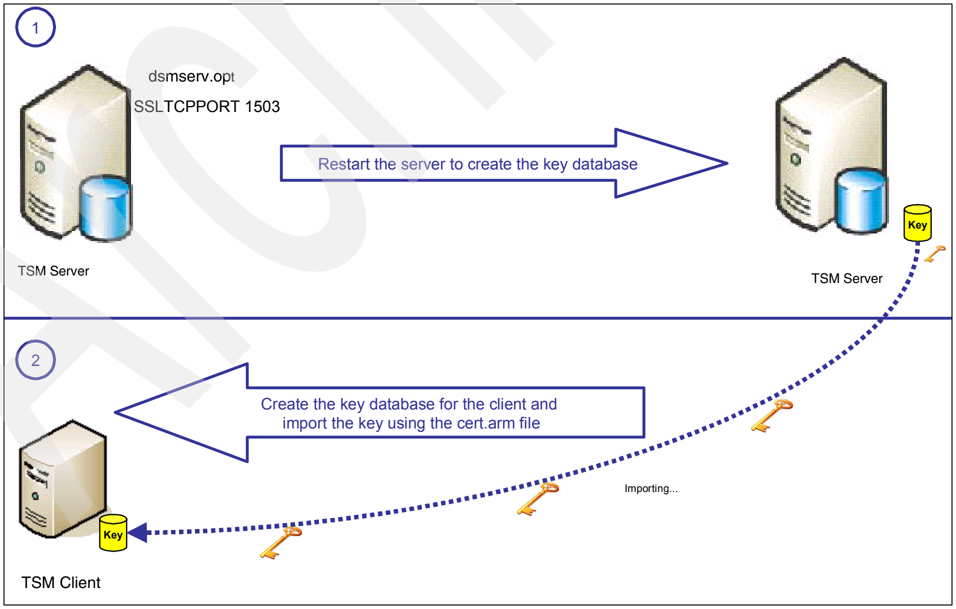

Figure 5-5 SSL configuration

Initially you perform server configuration, the new option SSLTCPPORT must be specified in the server's dsmserv.opt. After setting the option and restarting the server, a key ring database (cert.kdb) does not exist, and no password to access it is in the database. The key ring database is created. The following actions then take place:

1. The key ring database access password is automatically generated, and the encrypted password is stored in the Tivoli Storage Manager server database. This enables the server to open the key ring database and access the certificate information.
2. The Tivoli Storage Manager server generates a self-signed certificate and stores it in the key ring database in the server instance directory. The Tivoli Storage Manager server was started in this directory, and it stores the dsmserv.dsk file there.
3. The public certificate that can be used by the Tivoli Storage Manager client is extracted and put in the cert.arm file. This file must be imported to the client.

You then configure the client. The following client components support SSL:

- Command-line client
- Administrative command-line client
- Backup-archive client GUI
- Client API

Only outgoing client/server connections support SSL. Incoming connections (for example, CAD, server-initiated schedule connections), do not support SSL. Client-to-client communications and the Web GUI also do not support SSL.

You have to create the local key database as long as the server's certificate cert.arm is available. This process creates the following files:

dsmcert.kdb	The key database.
dsmcert.sth	Stash file to hold the stored password. The password is used to encrypt the key database; the stash file is used later by the Tivoli Storage Manager client to retrieve the key database password.

After creating the database, you need to copy the server certificate cert.arm from the server machine to the client. When you do so, ensure that the transfer is secure. The key is then imported using the GSKit from the file cert.arm file into the client's key database. Then you delete the local cert.arm file because it is no longer needed.

You are ready to enable SSL communication between the Tivoli Storage Manager client and server; simply add the SSL YES option in the client options file. The value of the TCPPORT parameter (see “Client communication processing” on page 215) must match the server’s SSLTCPPORT value.

Figure 5-6 shows the SSL communication process after the initial configuration is performed as previously described.

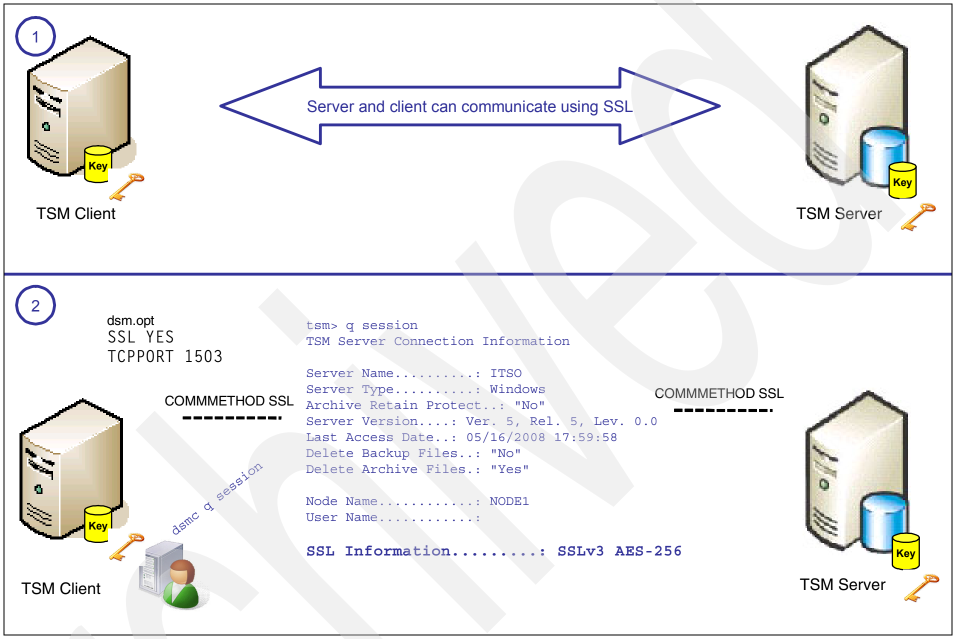

Figure 5-6 SSL secure layer communication

SSL adds another level of security to your implementation, and this level of security is easy to configure. It effectively secures client/server authentication and provides encryption for the rest of the session. Only the SSL communication layer is encrypted. The Tivoli Storage Manager client passes the data on the host as plain text; then SSL encrypts the data, and the data is sent over the wire to be decrypted by the SSL layer on the server end before being stored in the Tivoli Storage Manager server. This process means that the client and server are not aware that the data is encrypted; that is, the encryption is transparent to them.

If the client wants to secure the communications path between the client and the Tivoli Storage Manager server but does not want (or need) to store the data on the server in encrypted form, SSL is a better choice. SSL communication is recommended for remote connections to the server, for example, over the Internet.

Figure 5-7 shows how SSL communication works.

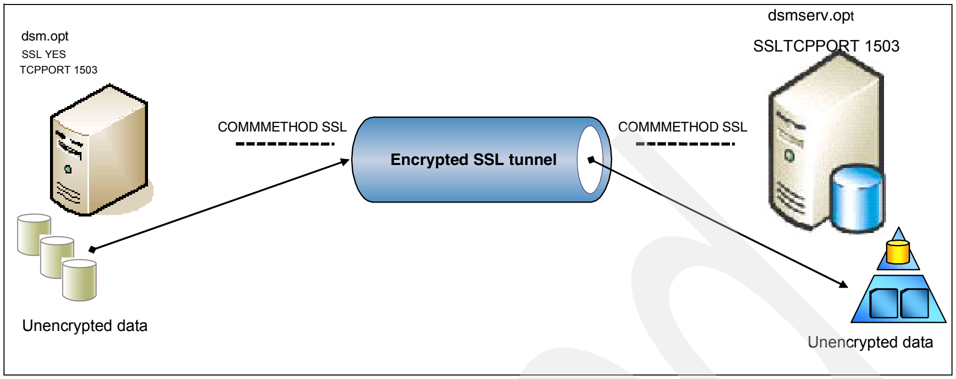

Figure 5-7 SSL communication layer encryption

For more information about SSL configuration and recovery options, refer to Chapter 21, "Version 5.5 client supported environments," *IBM Tivoli Storage Manager Versions 5.4 and 5.5 Technical Guide*, SG24-7447.

5.2.7 IBM Tivoli Storage Manager tape hardware encryption

IBM offers hardware encryption on the IBM System Storage TS1120, TS1040, and LTO4 tape drives. Tape hardware encryption performs AES256 encryption when data is written on the device. This encryption is performed without any loss of performance when software-based encryption methods is used.

The implementation of IBM TS1120 is described in *IBM Tivoli Storage Manager: Building a Secure Environment*, SG24-7505. The implementation of IBM TS1040 is described in *IBM Tivoli Storage Manager Versions 5.4 and 5.5 Technical Guide*, SG24-7447.

Data encryption can be enabled or managed at one of three levels: application, system, or library, as shown in Figure 5-8.

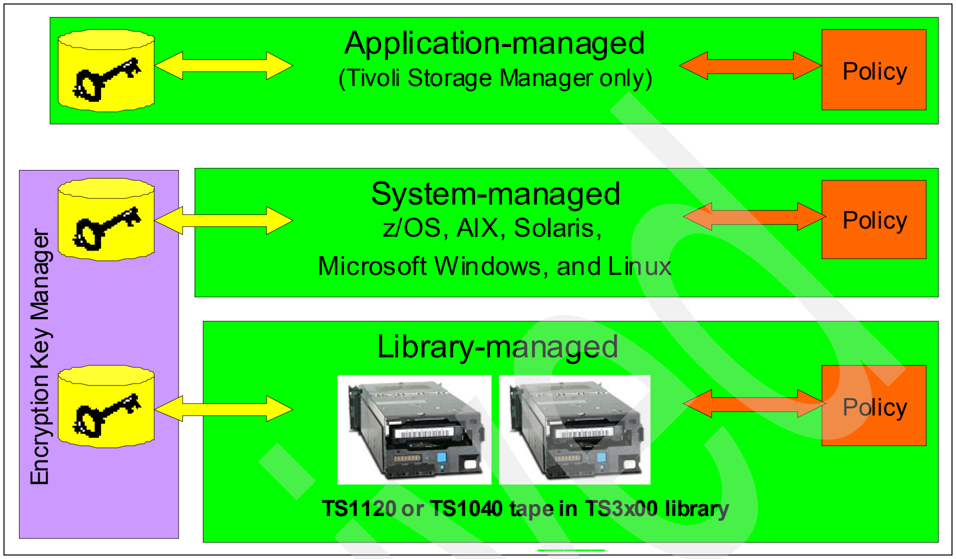

Figure 5-8 Tape encryption methods

The following methods enable tape encryption:

- Application-managed encryption (AME) is managed within a software application that uses the tape drive. Currently, IBM Tivoli Storage Manager is the only application that supports IBM tape encryption. In this case, Tivoli Storage Manager manages the keys that are used for encryption.
- System-managed encryption (SME)
- Library-managed encryption (LME)

The latter two methods, SME and LME, use an external, separately installed IBM Encryption Key Manager (EKM).

The methods have the following differences:

- Where the encryption policies reside
- Where key management is performed
- Whether an EKM is required
- How the tape drives communicate with EKM

Implemented in the tape drive, data is encrypted before it is written to the cartridge. When tape compression is enabled, the tape drive first compresses the data to be written and then encrypts it. This process means that IBM tape encryption involves no loss of capacity. If the encryption solution encrypts the data first and then tries to compress it, the encrypted data usually compresses very little if at all.

Important and sensitive data can be protected in many ways. Data can be encrypted by means of special software programs, hardware adapters, facilities, or outside of the device where the data is stored. Encrypting data with software programs takes away processor power, and encrypting data with hardware requires additional investment in hardware for the computers.

Encrypting in the drive also eliminates the need for any additional machines or appliances in the environment by offloading the encryption processing overhead onto the drive. Because the drive can also process unencrypted workloads, the IT environment is further simplified, eliminating the need for separate drives to process data that does not have to be encrypted.

The IBM tape encryption solution uses a combination of symmetric and asymmetric encryption methods. For more information, refer to 5.2.1, “Introduction to encryption” on page 211. The IBM tape encryption solution uses a symmetric AES data key to encrypt and decrypt data. This data key is protected by the asymmetric RSA algorithm and is not available in the clear when tape drives and the Enterprise Key Manager (EKM) communicate.

The sole task of the EKM is to handle key management. The EKM does not perform any cryptographic operations, such as generating encryption keys.

For application-managed encryption (AME), the application has to be capable of generating and managing encryption keys and of managing encryption policies. At the time of writing, the only application with this capability is Tivoli Storage Manager. Policies specifying when encryption is to be used are defined through the application interface. The policies and keys pass through the data path between the application layer and the encrypting tape drives.

Encryption is the result of interaction between the application and the encryption-enabled tape drive and does not require any changes to the system and library layers.

AME is the easiest encryption method to implement and adds the fewest responsibilities for the storage administrator. Because the data path and the key path are the same, data and drive availability are not at risk. Policy granularity depends on the application. With Tivoli Storage Manager, you control encryption on a storage pool basis. AME does not use centralized key management because the application generates, stores, and manages the encryption keys. The lack of centralized key management makes tape interchange and migration more difficult.

AME can be the most convenient solution when Tivoli Storage Manager is the only application that utilizes tape encryption.Tivoli Storage Manager does not restrict you to using AME. You can also choose SME or LME to encrypt Tivoli Storage Manager data.

Figure 5-9 depicts Tivoli Storage Manager tape encryption. The text after the figure uses the numbers in the figure to describe the steps in the process.

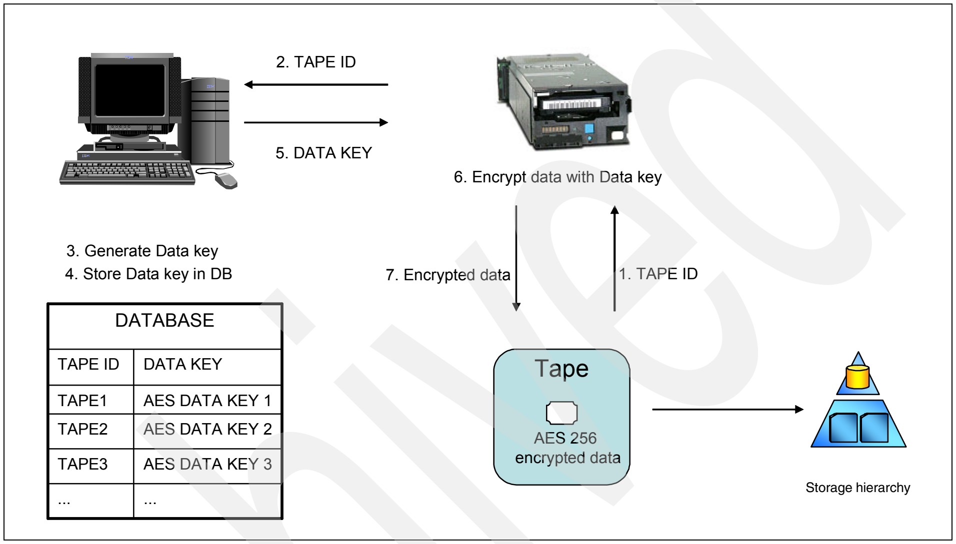

Figure 5-9 Tivoli Storage Manager tape encryption

The following steps describe how data is encrypted to tape using Tivoli Storage Manager as the key manager:

1. The tape drive mounts a tape for encryption.
2. The tape drive sends its TAPE ID or VOLSER to Tivoli Storage Manager.
3. Tivoli Storage Manager generates a 256-bit AES data key and encrypts the data key.
4. Tivoli Storage Manager stores the encrypted data key and the tape identifier in its database.
5. Tivoli Storage Manager sends the data key to the tape drive.
6. Tape drive uses the AES algorithms to encrypt data
7. Encrypted data is written to tape.

For more information about tape encryption refer to the books:

- *IBM System Storage Tape Enterprise Key Manager, Introduction Planning and User Guide*, GA76-0418.
- *IBM Tivoli Storage Manager: Building a Secure Environment*, SG24-7505.

- *IBM System Storage Tape Encryption Solutions*, SG24-7320.
- *IBM Tivoli Storage Manager Versions 5.4 and 5.5 Technical Guide*, SG24-7447

Implementing tape encryption secures your data when your tapes leave the facility. It is extremely important if your environment requires tape transportation, either manually by operators or externally, sending it to a backup site or another storage facility.

Determine your business continuity plan and make sure that your backup data is secure and duplicated for disaster recovery purposes. As administrator, you need to be aware of the library inventory. Tivoli Storage Manager helps identify whether your tapes are in or out of your tape library, and what the status is of each volume. The communication between the tape operators and you, as the Tivoli Storage Manager administrator, must be constant, minimizing the risk and improving your tape manageability.

5.2.8 Recommendations for enhanced security

The Education Company requirements are addressed by the following Tivoli Storage Manager security features:

- Provides better control of regular user's access to the Tivoli Storage Manager server.
- The administrators created in the Tivoli Storage Manager server match the user names.
- Applies policy rules for password credentials.
- Implements client-side encryption for communication between the file servers, application, and database servers.
- Implements SSL communications for mobile users, including instructors who connect to the network and back up their files over the Internet.
- Establishes Tivoli Storage Manager server in a firewalled environment and sets communication ports to different values than the default.
- Provides better control and management of tapes, enabling tape encryption to ensure confidentiality during data transport.

Security options and recommendations are illustrated in Figure 5-10.

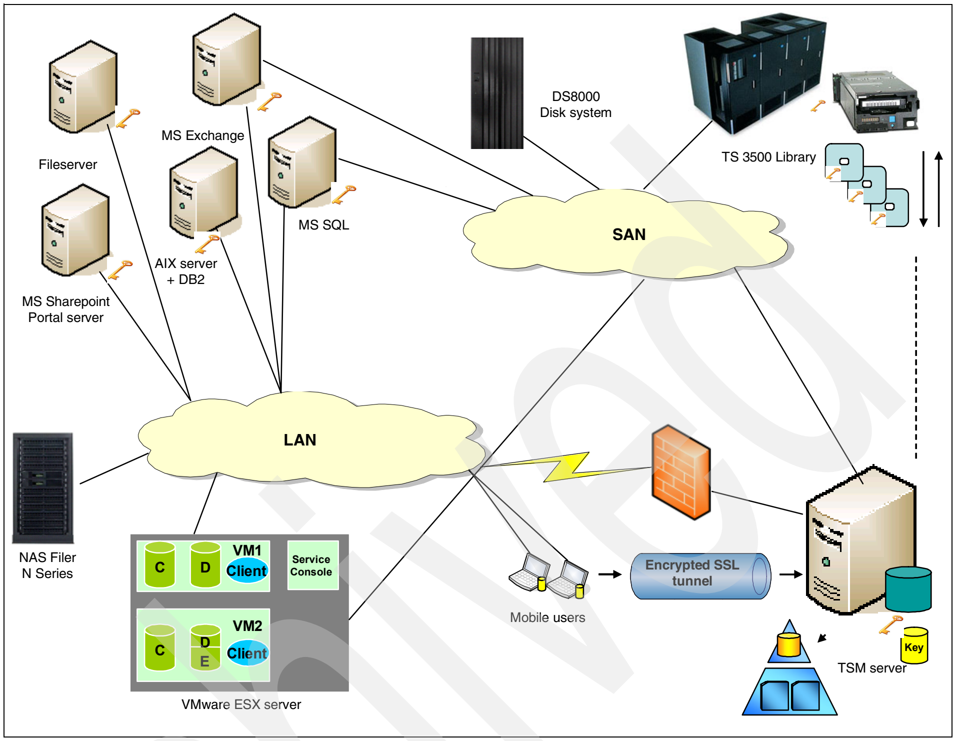

Figure 5-10 Recommended security environment at the Education Company

5.3 Network Data Management Protocol (NDMP)

Network Data Management Protocol (NDMP) is an open standard network protocol that enables storage-management applications to control backup and recovery of an NDMP-compliant Network Attached Storage (NAS) device over the network. This process avoids installation of third-party backup software on the protected device. An NAS vendor deals with file systems, bulk data transfer, and secondary storage devices. The backup vendors can then focus on control of backup and recovery operations and on data management.

With Tivoli Storage Manager Extended Edition and the NDMP, we can perform high-performance, scalable backups and restores of an NAS file server. Tivoli Storage Manager Extended Edition and the NDMP provide full or differential backup of a file system image at a directory or file system level. Tivoli Storage Manager Extended Edition and the NDMP also allow restores of an entire file system or selected files and directories within the file system of Network

Appliance file servers with OS Data ONTAP V6.1.1 or higher and EMC Celerra systems running Dart V5.1.9.3 or higher. Multiple backup and restore operations can be performed simultaneously. Data ONTAP supports a maximum of 16 concurrent backup and restore operations.

NDMP supports a maximum of 128 concurrent sessions on NearStore® systems and 40 concurrent sessions on other systems. General NDMP support also allows other NAS vendors to certify integration with Tivoli Storage Manager.

The NDMP backup and restore features are fully integrated with Tivoli Storage Manager Extended Edition server and client. No extra software is required on the server, client, or NAS appliance. When doing backups and restores, the NAS device and the Tivoli Storage Manager server and client all have specific roles, as shown in Figure 5-11.

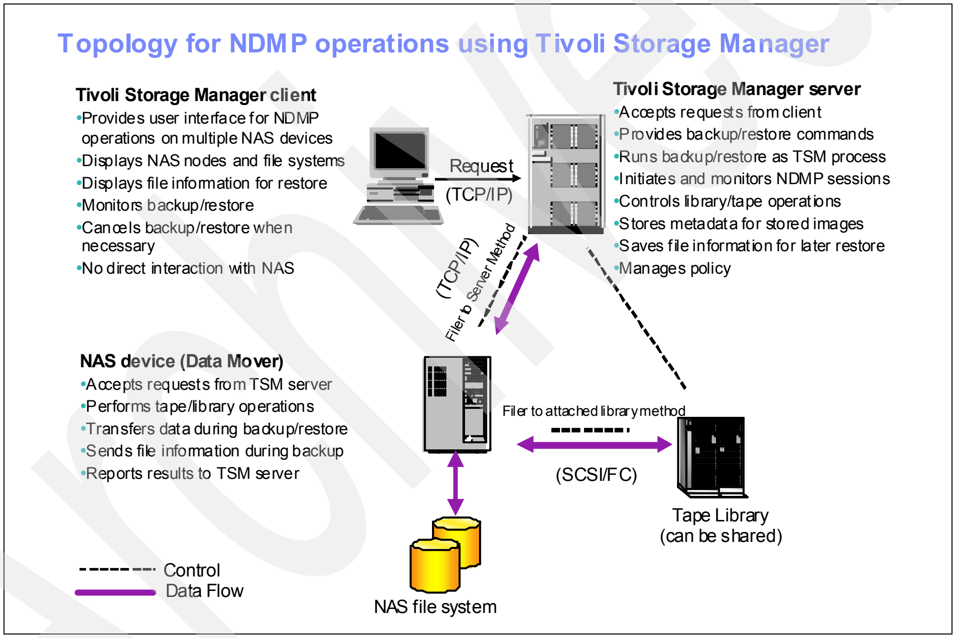

Figure 5-11 Topology for NDMP with filer to server and filer to attached library methods

Tivoli Storage Manager Extended Edition enables file-level and full/differential file system image backups and restores of servers that support the NDMP protocol. During backup and restore operations, data flows directly between the tape drive and the NAS appliance. Alternatively, you can back up directly to the Tivoli Storage Manager storage hierarchy and also implement Disaster Recovery Manager (DRM) because it now supports NAS storage. Multiple backup and restore operations can be performed in parallel.

NDMP for NAS backup uses either a SCSI-attached tape device local to the NAS appliance, a SAN-attached SCSI, or an Automated Cartridge System Library Software (ACSLS) device that can be shared with the Tivoli Storage Manager server. Library robotics can be controlled directly by the Tivoli Storage Manager server or by passing SCSI commands to an NAS file server.

Drives must be supported by both the NAS appliance and the NAS OS. Drives can be dedicated to NDMP operations from a single NAS file server or can be shared. Multiple NAS appliances can share SAN-attached shared tape resources if backups are performed using the same Tivoli Storage Manager server. Drives can be also shared with LAN-free backup and restore operations, provided that the library is controlled directly by the Tivoli Storage Manager server.

5.3.1 Current environment backup methodology

The Education Company backs up the NAS filer today using the CIFS and NFS mapped drives technology. During backups, if the Education Company experiences memory problems and the backup takes a long time to complete, you sometime have to cancel the back up.

The NAS has access to the data and needs to make the data available though the network where the mapped drive is located. The current backup methodology is inefficient because the machine where the backup is taken requires locally mapped drives. Then once the backup is initiated, more resources are required to complete the backup in terms of network and capacity.

Tivoli Storage Manager was selected to address this issue. By enabling NDMP backups, Tivoli Storage Manager enables the client to move the data over the same network to the Tivoli Storage Manager server. This process addressees the network overload and resource issue.

5.3.2 Planning for NDMP

The critical part of planning for NDMP implementation is deciding where you plan to attach your devices and which devices you plan to use them for. You need to determine whether the tape library is attached directly to the NAS filer or whether the tape library is attached and controlled by the Tivoli Storage Manager server. Wherever the library is physically attached, the NAS filer must have access to the tape drives.

The distance between the NAS filer and the Tivoli Storage Manager server is a crucial factor to consider when deciding whether the library must be attached to the NAS. Once you determine library connectivity, you have to decide whether to dedicate the library solely for NDMP operations or share the tape library and drives with the Tivoli Storage Manager server.

It is important to know what to do with the data after backing up to the Tivoli Storage Manager. If you want to perform any data movement processes, the method of implementation is crucial.

For back-end data movement, introduced in Version 5.4, you need at least two mount points or physical paths (or both), available for a single NAS. You can mix these paths in different types of media and libraries.

At this point, you can then answer the following questions:

- Which type of library (SCSI, ACSLS, or 349X) is going to be used?
- If it is a SCSI library, where is it going to be attached?
- Do you want the data to go directly to tape or use disk?
- How do you want to use the tape drives? Are the tape drives going to be dedicated to Tivoli Storage Manager or to NAS, or are they going to be shared between NDMP and Tivoli Storage Manager operations?
- Is the data going to be backed up tape-to-tape for disaster recovery functions?
- Do you plan to send the backup data to a single Tivoli Storage Manager server instead of attaching a tape library to each NAS device?
- Do you want to keep all hardware on the Tivoli Storage Manager server and send NDMP data over the LAN?

Based on the answers to the preceding questions and the information provided earlier in this section, you can now decide which method to implement.

Tivoli Storage Manager server Version 5.4 and higher supports NDMP V4 in two different configurations:

- Filer to attached library
- Filer to server (new in V5.4)

For more information about the NDMP prerequisites, see “NDMP support using IBM Tivoli Storage Manager” on page 15.

Filer to attached library

In the filer-to-attached-library configuration, the NAS device has access to a locally attached tape library. Either the NAS device or the data management application, which is the Tivoli Storage Manager server, controls the library robotics with SCSI or Fibre Channel. Either way, only the drives have to be attached to the NAS.

Filer to attached library is similar to LAN-free backups of NAS boxes. The NAS box writes the backups to tape and must be in the correct DUMP format for that specific box. The NAS box needs Fibre-Channel or SCSI connectivity to the tape device or library.

Using this method, NDMP backups are not part of the Tivoli Storage Manager storage hierarchy, because NDMP backups are not in a Tivoli Storage Manager native data format but use NAS format (NDMP dump images). Once the data is written in NAS format, the data movement is supported only between storage pools of the same dump format. Data cannot be moved between NDMP storage pool format and NATIVE storage pools. For this reason, it is not possible to use migration and reclamation tasks. The volume is reused once the image is expired by the NAS filer. The NAS filer writes the data to the volumes and manages the volumes.

If the library is directly attached to the NAS, NAS passes the control commands to the tape library. Both the Tivoli Storage Manager server and the NAS must be in the zone to see the library and the drives.

You can share the library with the Tivoli Storage Manager and the NAS box, or you can have the NAS box itself control the library. The difference is in how you define the PATH, as SRCTYPE=SERVER or SRCTYPE=DATAMOVER.

Figure 5-12 shows the filer-to-attached-library configuration.

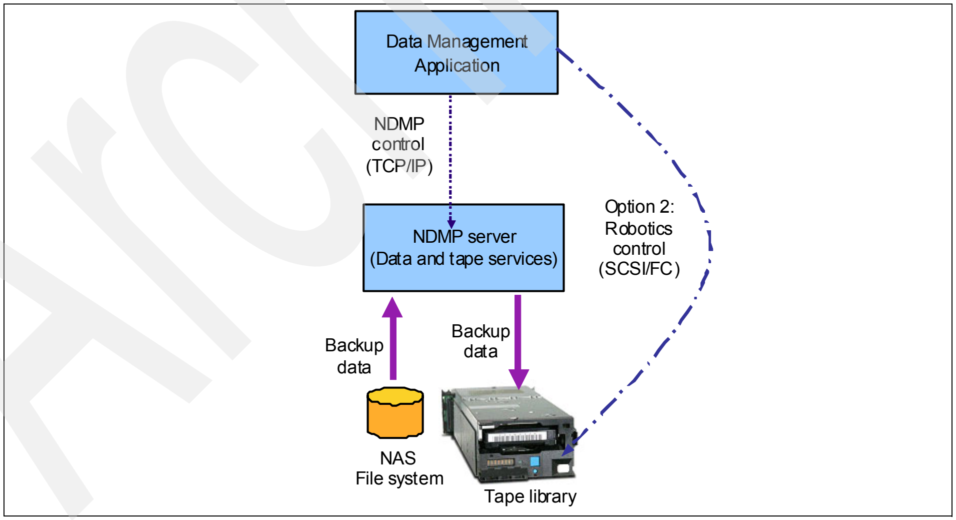

Figure 5-12 Filer-to-attached-library configuration

Filer to server

In the filer-to-server configuration, the library is attached to the Tivoli Storage Manager server. The NAS device does not have access to the library.

This configuration is also known as *three-way NDMP backup*. The backup data from the NAS device is transferred over the network (TCP/IP) to the Tivoli Storage Manager server.

Think of filer to server as LAN-based backups of NAS boxes. Data is sent over the network and stored in native storage pools like any other data. Tivoli Storage Manager is the target of the NDMP backup.

Figure 5-13 depicts the filer-to-server configuration method.

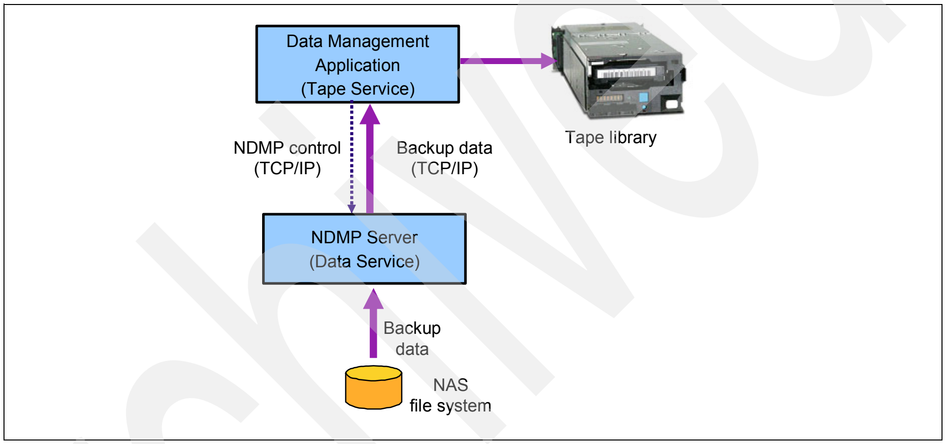

Figure 5-13 Filer-to-server configuration method

With the filer-to-server method, Tivoli Storage Manager writes the data, which enables you to move the data to another storage pool because the data is written in native data format, which is part of the Tivoli Storage Manager storage hierarchy.You can then perform migration and reclamation. The data is written in native format, as part of the TSM storage hierarchy.

5.3.3 Back-end data movement

As stated in the previous sections, Tivoli Storage Manager supports data movement between storage pools with the same data format.

Tivoli Storage Manager Version 5.4 and higher supports back-end data movement of tape volumes containing NDMP-generated backup images using the NDMP tape-to-tape copy function.

The Tivoli Storage Manager Server asks the NDMP tape-to-tape copy function to copy either the whole tape to another tape or only one or a few backup images from one tape to another tape within a single NAS device.

You can use the NDMP tape-to-tape copy function to create copy storage pools from primary storage pools that contain NDMP-generated backup images. This improves the data availability of the NAS file system and enables you to create volumes for off-site vaulting that can be managed by the Disaster Recovery Manager (DRM). Use the BACKUP STGPOOL command to do so.

In addition, you can use the MOVE DATA command to move the NDMP-generated backup images from one storage pool volume to another. The target storage pool must have the same NDMP data format as the source storage pool.

Data migration and space reclamation are not supported for storage pools using the NDMP data format. However, you can use intra-pool data movement (within the same storage pool) or inter-pool data movement (in different storage pools) for space recovery. You can perform intra-pool data movements with primary storage pools and copy storage pools. Use inter-pool data movement for migration to a new device type. Issue the MOVE DATA command to perform each data movement method.

To back up a primary storage pool that contains NDMP-generated backup images, you need a copy storage pool with the same data format as the primary storage pool on the Tivoli Storage Manager server.

For more information, refer to the Chapter 10, “NDMP,” *IBM Tivoli Storage Manager Versions 5.4 and 5.5 Technical Guide*, SG24-7447.

5.3.4 NDMP file-level and directory-level restore

During the backup, an image is sent to the Tivoli Storage Manager server. For restore purposes, you have to restore the entire image, but you have the option of creating a table of contents (TOC) of file information and storing it in the Tivoli Storage Manager storage pool. We recommend a disk storage pool for fast restores. You therefore have the following two options of restoring the files and directory structures:

- File-level restore
- Directory-level restore

NDMP file-level restore

NDMP file-level restore provides file-level restores from backup images produced by NDMP operations. To assist users in restoring selected files, you can create a table of contents (TOC) of file-level information for each backup image. Using the Web backup/archive client, users can then browse the TOC and select the files they want to restore. If you do not create a TOC, you must specify the name of the backup image that contains the file to be restored and provide the fully qualified name of the file.

You can create a TOC using one of the following commands:

- BACKUP NODE server command
- BACKUP NAS client command, with include.fs.nas specified in the client options file or in the client options set.

Using a TOC file, you can restore the objects individually and expand the directory tree as though it was a normal backup/archive client.

To do so, you need an additional client node in the Tivoli Storage Manager server and an administrator who administers both the NAS node created for backup and restore purposes and the node created for the Web client interface. You must explicitly specify both options, NASNODENAME and NODENAME, in the client options file.

You can configure the client acceptor daemon (CAD) to access the Web client interface. The Tivoli Storage Manager server loads the TOC file into its database, enabling the user to explore the NAS file systems. Figure 5-14 shows how Tivoli Storage Manager loads the TOC file.

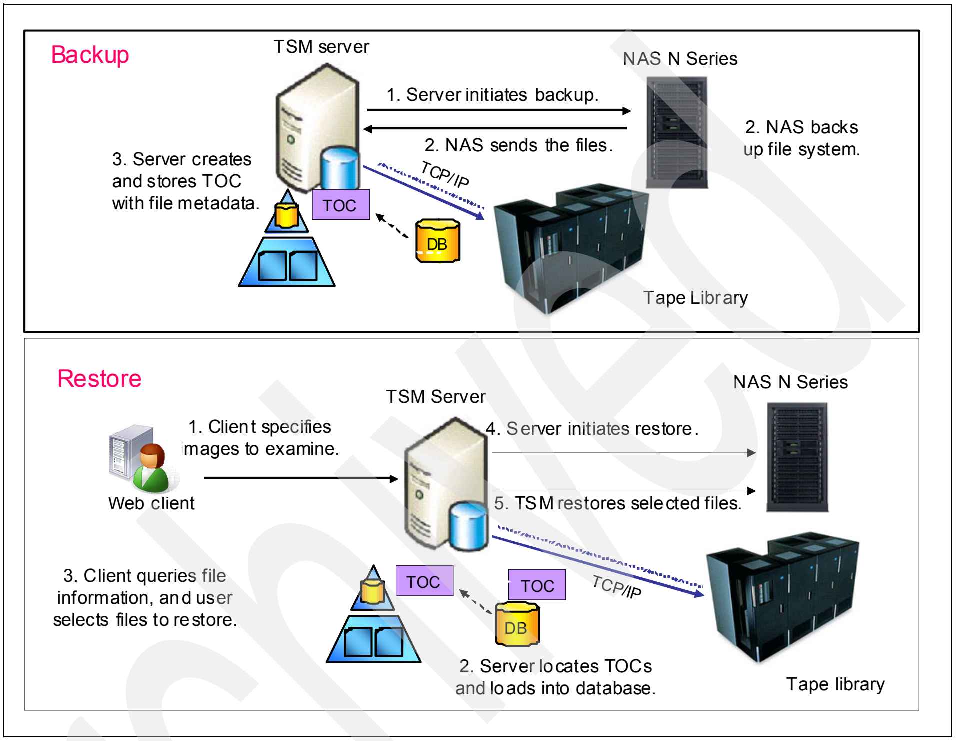

Figure 5-14 TOC file loaded into Tivoli Storage Manager database

Important: Remember you cannot restore a file directly on a shared mapped drive. The NAS device performs the data movement and must be able to access the target and restore the objects in the original location (such as the NAS file system). But if the restored object is in a mapped drive, it is available in the mapped drive once the restore is completed.

NDMP directory-level backups and restores

For a large NAS file system, initiating a backup on the directory level reduces backup and restore times and provides more flexibility in configuring your NAS backups. By defining virtual file spaces, you can partition a file system backup among several NDMP backup operations and multiple tape drives. You can also use different backup schedules to back up subtrees of a file system. The virtual file space name cannot be identical to any file system on the NAS node. If you create a file system on the NAS device with the same name as a virtual file system, a name conflict occurs on the Tivoli Storage Manager server when the new file space is backed up.

Using the Tivoli Storage Manager server DEFINE VIRTUALFSMAPPING command, you can define specific target points for backup and restores. For a large NAS volume, initiating a directory-level backup or restore reduces the total time of the operation.

When creating a virtual file space mapping, you have to specify the PATH of the target. The PATH in this case can be a folder in the file system specification. For example, the directory you defined in the PATH is the location of the objects you want to back up.

5.3.5 Recommendations for NDMP implementation

At the Education Company, the library is located near the Tivoli Storage Manager server. In this case, the filer-to-server method is the better choice for faster backup and restores. The company is planning to configure a corporate network, used only for backup and restore traffic.

The Education Company Tivoli Storage Manager server now holds a bunch of disks in the DS8000, which can be used for its disk storage pool. Because the data written to disk is native, it enables the Education Company to perform all Tivoli Storage Manager operations, such as move data, migration, reclamation, and to use the backup storage pool for disaster recovery purposes.

Figure 5-15 shows the filer-to-server configuration for the Education Company.

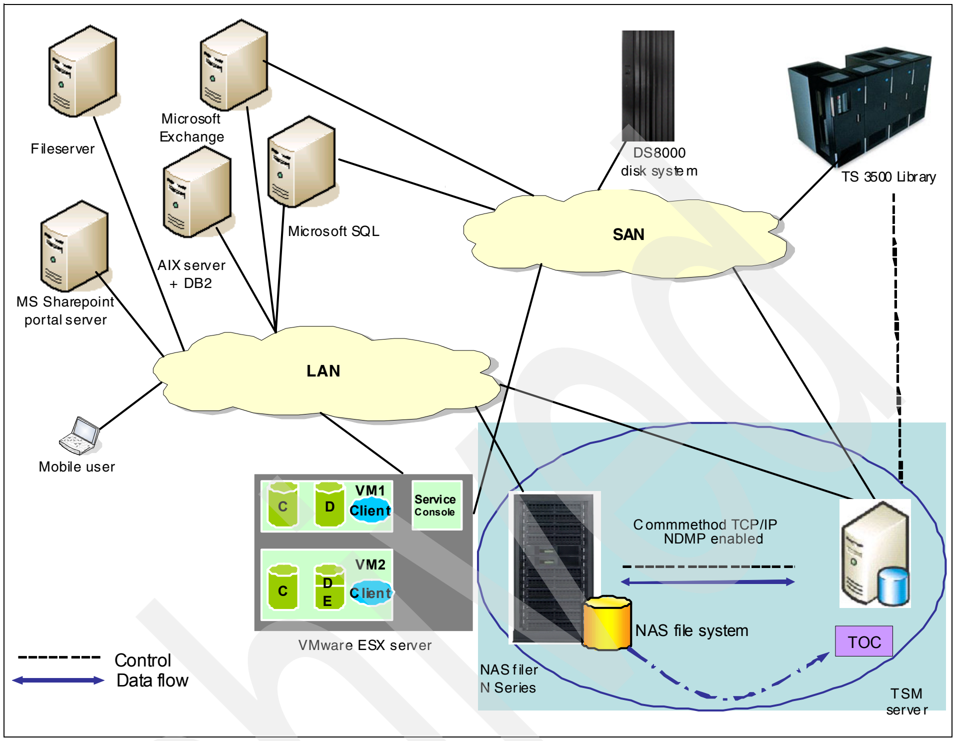

Figure 5-15 Filer-to-server implementation at the Education Company

With the NDMP implementation, the performance of the backup and restores is increased in addition to increased reliability of each image sent to the Tivoli Storage Manager.

5.4 Microsoft SQL database backup strategy

Information is the core of most organizations. Critical information is primarily stored in databases. Databases are where you can find an organization's intellectual property and information relating to its clients, suppliers, and other accounting details related to the business. To protect this data and to ensure the availability of the information stored on its databases requires a reliable backup solution.

Microsoft SQL Server is a relational database management system (RDBMS) produced by Microsoft. Its primary query language is Transact-SQL, an implementation of the ANSI/ISO standard Structured Query Language (SQL) used by both Microsoft and Sybase.

Databases Protection for Microsoft SQL Server is part of the Tivoli Storage Manager suite of complementary products. It uses the Tivoli Storage Manager for Databases for SQL API to communicate with the IBM Tivoli Storage Manager server, and it uses the Microsoft SQL Server API to communicate with the Microsoft SQL Server. This process is referred to as a legacy backup.

Note: We refer to the combined solution of the Data Protection for SQL Server together with Tivoli Storage Manager for Copy Services and the Tivoli Storage Manager backup/archive client generically as *Data Protection for Microsoft SQL Server*.

When you install IBM Tivoli Storage Manager for Copy Services together with Data Protection for SQL, you can also perform online snapshot backups to local shadow volumes, using Microsoft Volume Shadow Copy Services (VSS). These snapshot backups can be retained on the local shadow volumes and backed up to IBM Tivoli Storage Manager server storage. This process is also described as a *VSS backup*.

To reduce the overhead of backup operations on Microsoft SQL Server, you can choose to use another server to perform the backup to IBM Tivoli Storage Manager. Use a server with access to the shadow volumes, from either local snapshot volumes or from IBM Tivoli Storage Manager server storage.

For detailed information about backing up the Microsoft SQL Server databases, refer to *Backing up Microsoft SQL Server with IBM Tivoli Storage Manager*, SG24-6148.

5.4.1 Current Microsoft SQL server environment

The Education Company requires backing up its Microsoft SQL Server databases for both long-term and short-term retrieval. Each student in a class uses the databases as the primary learning tool. The class uses two types of databases. Individual students use a number of small databases that change daily; they also use larger databases that are not changed all the time.

The current backup strategy for the Education Company does not facilitate real-time backup and restore. Their strategy involves shutting down the database while it goes into “backup mode” to perform a full backup of the database. It is considered to be disruptive because database accessibility is interrupted during backup.

As a result, classes are disrupted to enable the restore of a database when required. The Education Company’s current method does not provide flexibility in terms of version management and data retention.

The Education Company must perform daily database snapshots on a seven-day period, so students can restore their individual databases at any given day within the week. This process enables the students to go back to a point-in-time to meet class requirements.

The original databases used in class must be kept for a longer period so they can be restored after a given classes’ changes are complete, in preparation for the next class. The Microsoft SQL Server database backup must be kept for 12 months, with monthly backups performed at the beginning of each month.

5.4.2 Planning for Microsoft SQL Server backups

A snapshot is a point-in-time copy of data on a disk volume. A snapshot copy typically is performed instantly and made available to other applications such as data protection, data replication, data analysis, and reporting. The original copy of the data continues to be available to the applications without interruption, while the snapshot copy is used to perform other functions on the data.

By utilizing the snapshot capability of VSS backups, each student benefits because the backup and recovery time is minimal. Online backups and restores are achieved, causing less disruption in class. Database backups can also be retained for a longer period of time by incorporating legacy backups through the Data Protection for Microsoft SQL Server application, with the Tivoli Storage Manager server managing backup policies.

Legacy backup

A legacy backup is a specialized API backup that functions with the Microsoft SQL Server storage engine, as shown in Figure 5-16. This type of backup is provided by previous releases of Data Protection for SQL.

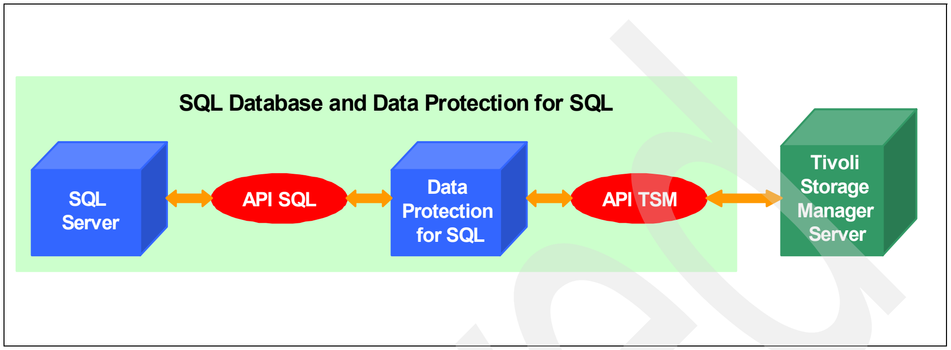

Figure 5-16 Data Protection for Microsoft SQL Server legacy backup communications

A legacy backup creates a copy of all or part of a Microsoft SQL Server database on Tivoli Storage Manager storage media. Data Protection for SQL provides selection mechanisms and the logic needed to back up and restore SQL data.

VSS backup

A VSS backup uses Microsoft VSS technology to produce an online snapshot (point-in-time consistent copy) of the Microsoft SQL Server database. The VSS backup can be stored on local shadow volumes or on Tivoli Storage Manager server storage.

A VSS backup means Microsoft SQL Server is not in “backup mode” for an extended period of time because the length of time required to perform the snapshot is usually measured in seconds and not hours. In addition, a VSS backup takes a snapshot of large amounts of data at the same time because the snapshot works at the volume level. VSS backups require the installation of the Tivoli Storage Manager for Copy Services, in addition the installation of Data Protection for SQL and the Tivoli Storage Manager backup/archive client. See Figure 5-17 on page 247 for a depiction of the general VSS architecture. You can optionally have an alternate machine move the data to the Tivoli Storage Manager server, which is an off-loaded backup.

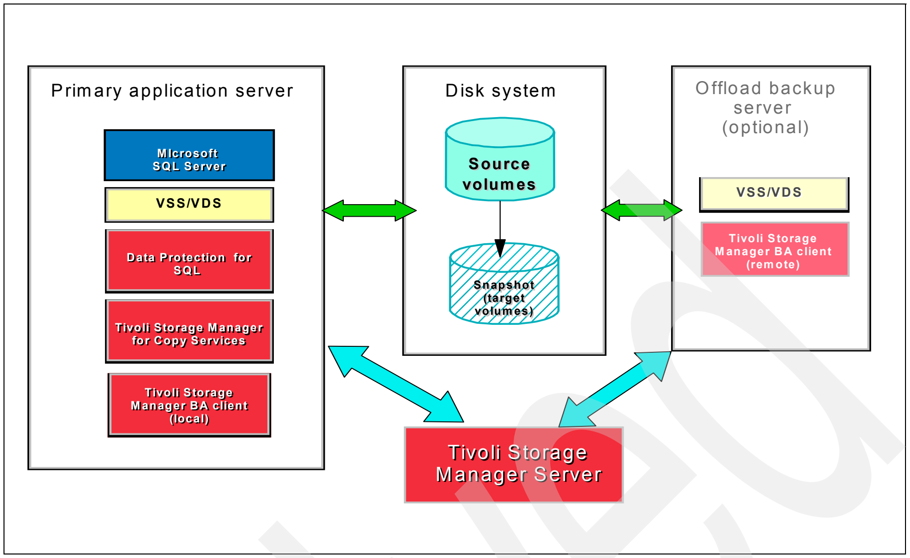

Figure 5-17 General VSS architecture

5.4.3 Combining VSS backup with legacy backup

Legacy backups enable you to perform full and incremental backups. A full legacy backup for a large database might require an extended period, such as several hours. Using only legacy backups, you are limited to performing a few full backups, such as weekly backups or even monthly backups. You can supplement full legacy backups by performing incremental backups multiple times a day to ensure database recovery.

Microsoft supports and recommends using both backup methods in your complete backup strategy. A mixture of a monthly legacy backup and a daily full VSS backup allows you a quicker recovery time. However, Microsoft also states that you cannot mix the two types of backups. For example, a legacy differential backup cannot be applied to a VSS full backup. Legacy differential and legacy log backups can be applied after a full VSS backup has been restored.

Best practices:

- The choice of legacy or VSS backups to Tivoli Storage Manager server storage is usually dictated by time, not versions.
- The number of backups to local shadow volumes are usually dictated by versions because of space limitations and VSS storage provisioning.

IBM Tivoli Storage Manager for Copy Services helps you back up and restore critical data so that it remains available 24x7. Using IBM Tivoli Storage Manager for Copy Services, you can implement high-efficiency backup and restore processes for your critical business applications with little impact on production performance. IBM Tivoli Storage Manager for Copy Services provides the following features:

- Enhanced backup and recovery features are integrated with existing Microsoft SQL Server backup and restore capabilities.
- Delivers fast backup with minimal impact on your production SQL server with a variety of snapshot providers.
- Provides near instant restoration of Microsoft SQL Server databases from a shadow copy image to the production volumes.
- Features policy-based management of multiple local backup versions so that recovery for multiple versions is easy.
- Improves availability and performance of production database servers with offloaded movement.
- Integrates snapshot capabilities with IBM Tivoli Storage Manager and its data protection component for Microsoft SQL Server.

5.4.4 Restoring SQL databases

When you restore a database, keep in mind that data that exists in the database is overwritten and is no longer available after the restore is complete. The command Restore Databases enables you to restore databases or parts of databases from full, differential, and log backups. Although only VSS full backups are supported, you can apply legacy differential and legacy log backups after you have restored a full VSS backup. You can restore SQL databases using the following restore methods:

- Legacy restore

  A Data Protection for SQL Server legacy restore obtains backup copies of all or part of one or more SQL databases and returns them to the SQL server. A complete restore of a database involves restoring a full backup or the equivalent thereof (from group, file, or set backups) and restoring all transaction logs after the last full backup.

- VSS restore

  A VSS restore restores VSS backups (SQL database files and log files) that reside on Tivoli Storage Manager server storage to their original location. The following characteristics are true for VSS restores:

  – VSS supports only full backup.

  – Restore granularity is at the database level.

  – Data must always be restored to the same drive letters and paths.

  – Restores can be performed on a Microsoft Cluster Server (MSCS) or Veritas Cluster server (VCS) environment.

- VSS fast restore

  A VSS fast restore restores VSS backups that reside on local shadow volumes. In general, restore processing can conclude within minutes instead of hours with a VSS fast restore. The following characteristics are true of VSS fast restores:

  – VSS supports only full backup.

  – Restore granularity is at the database level.

  – The key component of completing a VSS fast restore is the speed at which the application can become operational with the data that resides on local shadow volumes. Be aware that even though the data is restored relatively quickly, the transaction logs must still be replayed after the restore, and therefore increases the time of recovery for the application.

5.4.5 Recommendations for Microsoft SQL Server backups

To satisfy Education Company requirements, an SQL database backup is performed using a legacy backup that can be stored for a longer time. Data is stored for up to 12 months, with a monthly backup taken for historical purposes.

Real-time backups can be performed by way of snapshots where data retention is managed by versioning, and long-term storage can be achieved using the legacy backup methodology for a longer retention period. A daily snapshot is performed using VSS for up to five versions so the students can back up a one-week period, without losing too much information during class.

Legacy backups are stored on tape media utilizing tape storage pools for longer retention and historical data. VSS backups are stored on the SAN disks enabling a faster restore time.

Using VSS and legacy backups together can implement a highly effective backup solution. In the Education Company environment, the SQL database is backed up to the TS3500 library for the legacy backup. VSS backups are stored in the DS8000 disk system for fast restores.

In the Education Company environment, the DS8000 disk system utilizes the disk system purely for storage purposes. Refer to 5.4.7, “Optional VSS functionality using hardware devices snapshot” on page 251 for more information about this option.

The Microsoft SQL backups are processed LAN-free, using SAN. In the Education Company environment, the SQL DB itself is actually stored in the same DS8000 disk system. Figure 5-18 illustrates how the data is backed up.

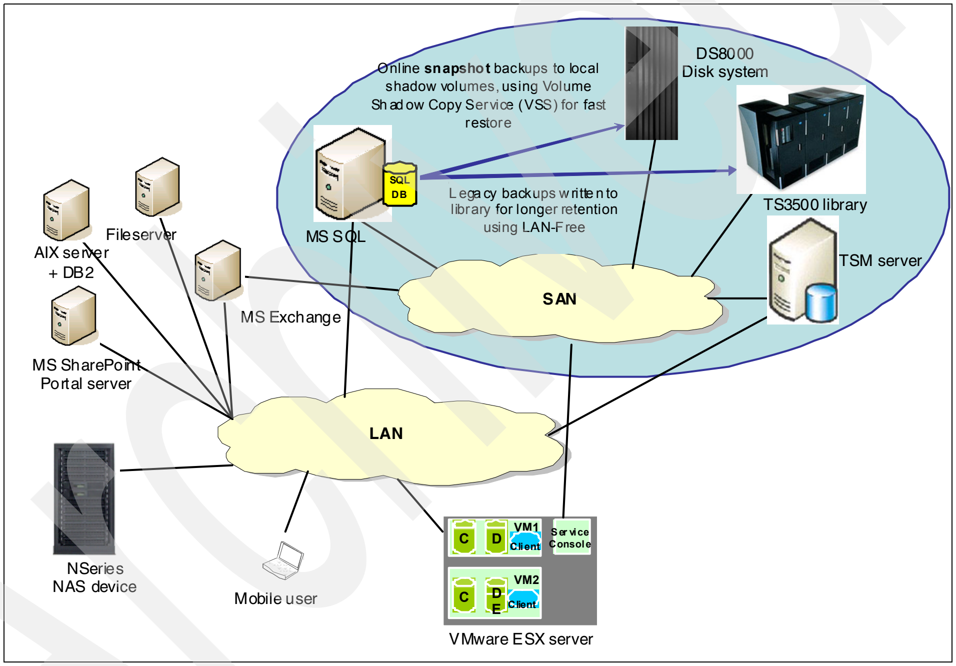

Figure 5-18 VSS backup architecture

5.4.6 Application requirements

For detailed installation and configuration information, refer to Chapter 3, "Installation and configuration," in *Backing up Microsoft SQL Server with IBM Tivoli Storage Manager*, SG24-6148. Data Protection for SQL requires the following applications:

- Tivoli Storage Manager backup/archive client Version 5.5.0 (or later)
- Tivoli Storage Manager API Version 5.5.0 (or later)
- Tivoli Storage Manager server Version 5.4.0 (or later)
- Microsoft SQL Server 2005
- IBM Tivoli Storage Manager for Copy Services

5.4.7 Optional VSS functionality using hardware devices snapshot

In some cases, it is beneficial to perform a snapshot on the hardware level for transportability. This topic is not covered in this book. It is introduced to provide additional information about how to back up the SQL database. However, you can find more information about this feature in *Backing up Microsoft SQL Server with IBM Tivoli Storage Manager*, SG24-6148.

You can create a transportable shadow copy that enables you to move the database from one server to another server. However, to facilitate this function, you have to install an additional module, the Hardware Devices Snapshot Integration module. Support for this feature is also limited to a number of storage subsystems.

The Hardware Devices Snapshot Integration module is used in conjunction with the Tivoli Storage Manager Microsoft Windows client. Together, the module and client implement the VSS Instant Restore functionality. For more information, refer to the requirements for the Tivoli Storage Manager V5.5 backup/archive Windows client.

VSS Instant Restores are only supported by the IBM System Storage SAN Volume Controller or the DS6000 and DS8000 storage subsystems in association with the VSS hardware provider. Therefore, you must install and configure IBM TotalStorage Support for Microsoft Virtual Disk and Volume Shadow Copy Services as your VSS hardware provider to perform VSS Instant Restores.

5.4.8 Snapshot backups for other applications

You can also utilize snapshot technology to back up other applications such as Microsoft Exchange and DB2 databases.

Microsoft Exchange database backup with VSS

Data Protection for Microsoft Exchange operations use the Tivoli Storage Manager API to communicate with the Tivoli Storage Manager server and use the Exchange API to communicate with the Exchange server. In addition to using these APIs, Data Protection for Microsoft Exchange VSS operations also use the Tivoli Storage Manager backup/archive client and Microsoft VSS technology to produce an online snapshot (point-in-time consistent copy) of Exchange data that can be stored on local shadow volumes or on Tivoli Storage Manager server storage. For more information, refer to the *IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange*, SC32-9058-04.

The Microsoft Exchange VSS Integration module is used in conjunction with Data Protection for Microsoft Exchange and the Tivoli Storage Manager Windows client. Together, these components implement the VSS requestor interface to drive the Microsoft Volume Shadow Copy Services for backup and recovery of the Exchange server.

With the addition of the Hardware Devices Snapshot Integration module, you can also get the VSS Instant Restore function on the SAN Volume Controller (SVC), DS6000, and DS8000. This component is simply a license file. When you install this component, it enables VSS operations within the Data Protection for Microsoft Exchange Server interfaces. To use VSS features, you must also install the module from IBM Tivoli Storage Manager for Copy Services:

Microsoft Exchange VSS
Hardware Devices snapshot

Refer to the requirements for Data Protection for Microsoft Exchange Server at:

http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/topic/com.ibm.itsmfm.doc/dpexc53.htm#dpe_ref_ins_hwreqs

To use VSS features, you must also install the module from IBM Tivoli Storage Manager for Copy Services:

- IBM Tivoli Storage Manager for Copy Services Microsoft Exchange VSS Integration module (required for basic VSS operations)
- IBM Tivoli Storage Manager for Copy Services Hardware Devices Snapshot Integration module (required for VSS Instant Restore operations)

DB2 database

IBM Tivoli Storage Manager for Advanced Copy Services uses copy services functions provided in the underlying storage hardware to perform the following tasks:

- Back up Oracle or DB2 databases to a local storage system or to a Tivoli Storage Manager server
- Restore databases from local storage or Tivoli Storage Manager storage

IBM Tivoli Storage Manager for Advanced Copy Services for DB2 minimizes the impact of performing backups of DB2 databases using Tivoli Storage Manager on database server systems. Tivoli Storage Manager for Advanced Copy Services for DB2 UDB offloads the transfer of backup data from a production database server to a backup server. The DB2 UDB database must reside on storage subsystem volumes. For detailed information, refer to *Data Protection for Snapshot Devices for DB2 Advanced Copy Services Installation and User's Guide Version 5.5*, SC33-8330-00.

Data Protection for Snapshot Devices supplements Tivoli Storage Manager Advanced Copy Services for DB2 software component provided with the DB2 High Availability feature as of DB2 Version 9.5. Data Protection for Snapshot Devices offers the following enhanced functionality compared to DB2 ACS:

- Offloaded transfer of backup data from the DB2 database server to Tivoli Storage Manager. The source for this transfer is a snapshot backup previously generated by DB2 ACS. The tape backup can be integrated with the request to generate the snapshot backup.
- Unlimited number of snapshot versions.
- Incremental and “no copy” FlashCopy variants (INCR, NOCOPY).
- Snapshot backup from a mirror set managed by the AIX LVM mirroring function (except N Series devices), and corresponding snapshot restore.

Consult the announcement letter or the *Installation and User Guide* for each product for additional details.

5.5 Microsoft SharePoint backup

In recent years, Microsoft Windows SharePoint Services and Microsoft Office SharePoint Server 2003 and 2007 have gained widespread popularity. This popularity is because of the products’ ability to easily facilitate team cooperation and efficiency. As a direct impact, Microsoft SharePoint is now demanding the same degree of availability and recoverability. This section focuses on Microsoft Office SharePoint Server (MOSS) 2007 capabilities and backup possibilities.

MOSS is a Web portal server that enables users to share information. It is predominantly used for sharing documents such as meeting minutes and contracts, and for sharing forms internally and externally on the Internet. Users from different departments can easily locate documents because the documents are indexed. This ease of use often makes MOSS the preferred place to store all kinds of files, and some organizations replace or complement their file servers with MOSSs. MOSS subsequently grows rapidly and easily reaches sizes in the TB range. This size poses a challenge for backup. To run smoothly, backing up 10 TB of data requires careful planning and execution.

A typical MOSS environment, shown in Figure 5-19, includes a Microsoft SQL Server (which holds all the SharePoint data), a Microsoft Office SharePoint Server 2007, and a number of front-end Web servers acting as the access point for users.

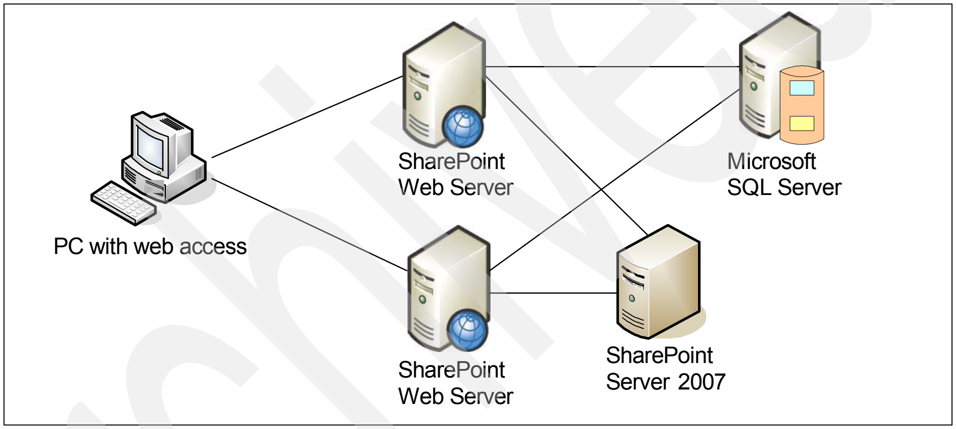

Figure 5-19 Typical MOSS environment

Figure 5-20 shows that each SQL database on the SQL server holds one or more MOSS site collections. These site collections are called a *site*, as in “Web site.” One or more subsites (or MOSS Web pages) are located below each site. These subsites contain the structure and actual files of each of the sites.

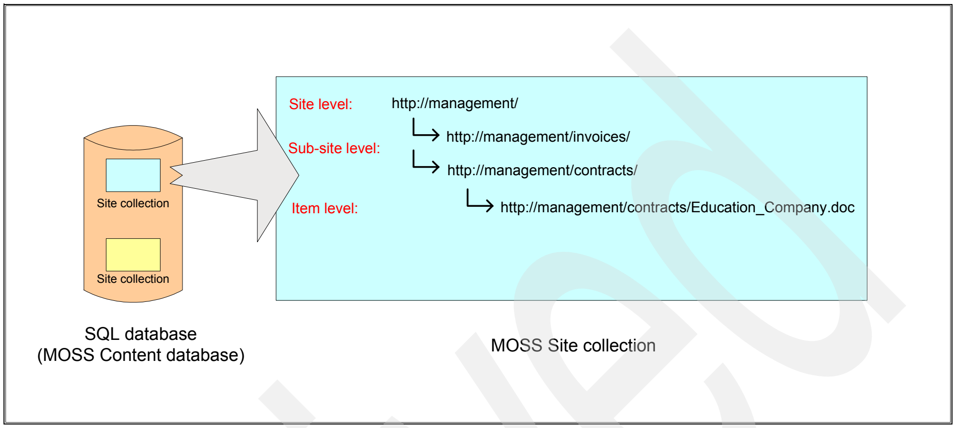

Figure 5-20 MOSS structure

5.5.1 Current MOSS environment

The Education Company has implemented a MOSS environment just as many other organizations have, but their current backup strategy does not provide the required functionality. The Education Company is facing increasing demands for flexible restores and recoverability. Their MOSS is growing exponentially and is currently 1.5 TB in size. The Education Company uses the following built-in backup and restore tools in MOSS:

Recycle Bin — A strategy (soft delete) designed to capture deletion events. Deletion events are defined as the files deleted in the same operation. A deletion event can be a single file or a whole document library full of files. The Recycle Bin enables some retention of content before the files are discarded, with a default retention of 30 days.

Central Backup — A native tool providing backup and restore at the MOSS farm level (multiple MOSS installations) or content database level (SQL database level).

This strategy has the following disadvantages:

- The Recycle Bin only handles deletion events. If an entire document library is deleted, single files cannot be restored.
- The Recycle Bin does not address data corruption. Only deleted files are backed up.
- The Recycle Bin works on a site-collection level. All deleted files have the same retention.
- It is not possible to perform on-demand backup and restore on the site, subsite, or file level. Only content databases (one or more sites) can be backed up.
- The Central backup only backs up to disk.
- The Central backup has no native scheduling options; only manual backup is possible.

Table 5-1 summarizes the current backup and restore possibilities.

Table 5-1 Current backup and restore strategy

Object type	Backup possibilities	Restore possibilities
Individual files and items	30 days soft delete	Files/items would have to have been soft-deleted before files can be restored. Only a whole deletion event can be restored.
Sites	Not possible	Not possible.
Subsites	Not possible	Not possible.
SQL databases	Central backup	Central backup.

The Education Company wants to enhance their backup and restore strategy and has the following requirements:

- Single file on-demand backup and restore
- Protection against data corruption
- On tape backup data storage
- Backups and restores scheduling
- Site- and subsite-level backups
- Backup jobs monitoring

5.5.2 Recommended MOSS environment

The Education Company chooses to implement Tivoli Storage Manager for Microsoft SharePoint because this complementary product for Tivoli Storage Manager can deliver the functionality the company requires.

Tivoli Storage Manager for SharePoint is a data protection module that provides centralized and automated backup and restore for single or multiple MOSS environments. The product is licensed from AvePoint, Inc., is integrated with Tivoli Storage Manager, and is able to store backup data in the Tivoli Storage Manager storage hierarchy. Due to the licensing, the individual Tivoli Storage Manager for SharePoint module is also referred to as the *DocAve module*.

New components

Our recommendation for the Education Company's implementation of the Tivoli Storage Manager for SharePoint architecture is shown in Figure 5-21 on page 258. The following servers and Tivoli Storage Manager for SharePoint modules are implemented:

- A server running Tivoli Storage Manager for Microsoft SharePoint Server (DocAve server)
- A server running Tivoli Storage Manager for Microsoft SharePoint Media Server (DocAve media server) and the Tivoli Storage Manager API
- One of the two SharePoint front-end Web servers installed with the Tivoli Storage Manager for SharePoint client (DocAve client)

To simplify the illustration, the added Tivoli Storage Manager for SharePoint components and servers are not shown with their proper connections to the LAN and SAN in Figure 5-21 on page 258. In reality, the SQL Server is connected to the SAN and all the servers are connected directly to the LAN.

For illustration purposes, all the Tivoli Storage Manager for SharePoint modules (server, media server, and client) have been separated and installed in separate server hardware in Figure 5-21 on page 258. These components can all coexist on the same physical server if needed or, as in the figure, be spread across multiple servers for better performance. The media server and client require somewhat powerful machines (depending on the size of the SQL database) because they are moving the actual backup data. The server component has a very small footprint and can easily be installed on an existing server. The Education Company has chosen to separate all the components for better performance and fault tolerance.

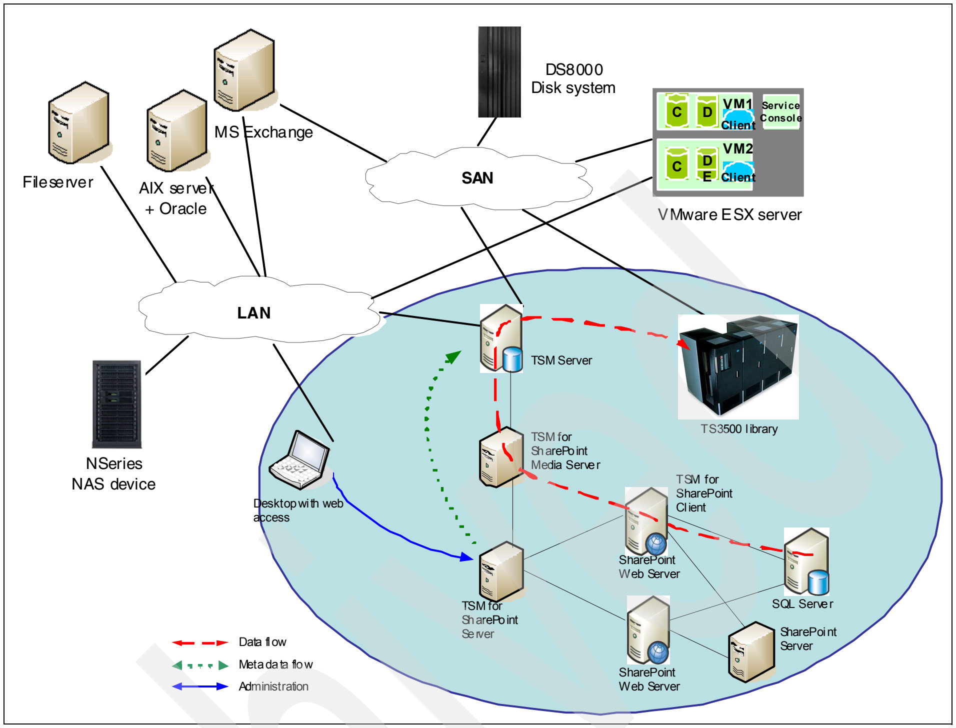

Figure 5-21 Recommended Tivoli Storage Manager for SharePoint architecture

Component features

Let us take a look at the role of the three components (server, media server, and client):

- The Tivoli Storage Manager for SharePoint Server is installed on this server, and it handles all configuration and administration of backup schedules, retention, job monitoring, include and exclude lists, and reporting. No backup data passes through this server. It acts purely as a configuration and administration interface for the rest of the Tivoli Storage Manager for SharePoint modules. The Tivoli Storage Manager Server functions *only* as a data repository.

  The Tivoli Storage Manager for SharePoint Server has no knowledge of what Tivoli Storage Manager for SharePoint is backing up or restoring. Only the Tivoli Storage Manager for SharePoint Server has knowledge of this and writes it to its own database, which is 10-20 MB in size. You can protect the database by backing it up to the local drive and afterwards backing it up to the Tivoli Storage Manager server using a special Tivoli Storage Manager backup/archive client (a client that just includes the Tivoli Storage Manager API). You can log onto the Tivoli Storage Manager for SharePoint Server administration interface through a browser, and you can use both native and Active Directory logons. This process provides effective security.

- The Tivoli Storage Manager for SharePoint Media Server is also installed. At the same time, a Tivoli Storage Manager API client is automatically installed side-by-side with the Tivoli Storage Manager for SharePoint Media Server. The Tivoli Storage Manager API client is used for the actual movement of backup data between the Tivoli Storage Manager for SharePoint Media Server and the Tivoli Storage Manager Server. The Tivoli Storage Manager for SharePoint Media Server moves the backup data between the Tivoli Storage Manager API client and the Tivoli Storage Manager for SharePoint Client.
- Finally, the Tivoli Storage Manager for SharePoint Client is installed on one of the MOSS front-end Web servers. This client extracts MOSS data directly from the SQL Server and sends it to the Tivoli Storage Manager for SharePoint Media Server.

Figure 5-21 on page 258 also depicts the flow of backup (and restore) data. When a backup is initialized by the Tivoli Storage Manager for SharePoint Server, the Tivoli Storage Manager for SharePoint Client begins to extract objects from the SQL database. Data is temporarily stored on the client until the backup finishes, and data is moved through the Tivoli Storage Manager for SharePoint Media Server (and Tivoli Storage Manager API client) and finally stored on the Tivoli Storage Manager Server. The metadata (the information about the objects being backed up) is sent directly from the Tivoli Storage Manager for SharePoint Server to the Tivoli Storage Manager Server. As mentioned, this metadata is not a complete list of objects being backed up. It is just sufficient information to enable the Tivoli Storage Manager Server to store the backup data on behalf of the Tivoli Storage Manager for SharePoint Server.

The data flow is of course reversed when doing restores.

Implementing Tivoli Storage Manager for Microsoft SharePoint

The Education Company gains the following benefits by implementing Tivoli Storage Manager for Microsoft SharePoint:

- Backing up data:
  - Integration with Tivoli Storage Manager and the ability to store MOSS backups in the Tivoli Storage Manager hierarchy, for example, on tape.
  - It is now possible for the Education Company to back up on the item, subsite, and site level, which provides a great deal of flexibility and enhances data security.
  - Backups on all levels can be run as full, differential, or incremental backups.

  - Individual file backup means individual retention on backed up files.
  - The company can perform on-demand backup, which ensures protection against data corruption.
  - Backup scheduling for multiple MOSS environments is centralized.
  - The company can perform custom backups. Schedules can back up a particular site each day, four subsites every six hours, or all the contracts of the company every hour (or all back up all three), so nothing is lost. Only imagination and storage space set the limit. You can exclude certain files from backup, so they are never backed up, thus saving storage space on the Tivoli Storage Manager server.
  - The company can monitor backup jobs and notification through e-mail or Microsoft Operations Manager (MOM). Administrators are automatically notified and able to rerun failed backups.
  - All backups are logged verbosely for future reference.
  - Backups can be compressed or encrypted, or both, before being sent and stored on the Tivoli Storage Manager server. This feature saves storage space and provides greater security.
- Restoring data:
  - The company can restore data at the level the backup was taken. Site-level backups result in site-level restores, subsite backups in subsite restores, and so on. The notable exception to this rule is item-level backups; you can use item-level backups for item-level restores as well as subsite and site restores.
  - The company can perform restores out of place to another subsite, site, or a different MOSS installation altogether. This feature makes it easy to move sites between MOSS installations.
  - The company can use out-of-place restores to promote subsites to sites and demote sites to subsites, within the same MOSS installation or between installations.

Table 5-2 summarizes the new backup and restore possibilities when using Tivoli Storage Manager for Microsoft SharePoint.

Table 5-2 New backup and restore possibilities

Object type	New backup possibilities	New restore possibilities
Individual files and items	Item level	Item level
Subsites	Subsite level	Subsite and item level (if using an item-level backup)

Object type	New backup possibilities	New restore possibilities
Sites	Site level	Site and item level (if using an item-level backup)
SQL databases	Site level	Site level

Best practices

Table 5-3 summarizes best practices for backing up and restoring MOSS. The way you tailor your own backups depends greatly on the importance of your own MOSS installation as well as on the individual files within it. We recommend a mixture of all four kinds of backups. Analyze your backup and restore demands carefully.

Your choice of backup level is also dictated by the speed of the different types. Site-level backups and restores are twice as fast as item-level backups and restores. SQL database backup is by far the fastest way to back up and restore but also the least flexible. Furthermore, item-level backups are stored as multiple objects in your Tivoli Storage Manager Server database, thus filling the database more rapidly. Site- and subsite-level backups are stored as fewer objects.

> **Note:** In Table 5-3, we recommend Tivoli Storage Manager for Databases to back up SQL databases, although we do not discuss this topic further in this chapter.

Table 5-3 Best practices

Object type	Backup	Restore
Individual files and items	Soft delete and item level	Soft delete and item level
Subsites	Subsite level	Subsite level
Sites	Site level	Site level
SQL databases	Tivoli Storage Manager for Databases	Tivoli Storage Manager for Databases

Recommended backup and restore strategy

The Education Company uses the new backup and restore functionality of Tivoli Storage Manager for Microsoft SharePoint and Tivoli Storage Manager for Databases and combines them.

The Education Company chooses the following strategy:

- Fourteen-day soft delete, so the company can restore deleted files from within MOSS itself.
- A daily item-level backup of the company's most important MOSS data (40 GB). A full backup occurs each Friday, an incremental backup is performed during the rest of the week.
- A daily subsite-level backup of two very important subsites (75 GB). A full backup occurs each Saturday, differential backup is performed during the rest of the week (for fast restores).
- The company performs a daily site-level backup of the most important sites that must be restored individually (500 GB). The company performs a full backup each Sunday and an incremental during the rest of the week.
- The company performs an SQL database backup of all data (all databases, 1.5 TB). It performs a full backup each Sunday and differential backup during the rest of the week. That way all data is backed up and can be restored. If even a single site backup is lost, all sites residing within that database are restored if the backup was performed with an SQL database backup.

5.6 VMware Consolidated Backup

Virtualization of hardware resources is more popular than ever. Virtualized storage and virtualized servers are a part of almost every IT installation. The popularity is attributed to the growing interest in freedom from geographical and physical limitations. In addition, virtual servers save energy; more efficient use of physical hardware (CPU, RAM, storage) results in less use of power and cooling, so that administration costs also decline.

VMware, Inc., provides software virtualization support for Intel-based hardware. It allows the use of multiple operating system images and operating system types, hosted from the same physical hardware, which in turn generally allows greater overall utilization of that shared hardware. We refer to each operating system image virtualized in this way as a *guest*.

A number of types of VMware servers are available. In the rest of this chapter, we discuss the VMware ESX Server because it is the enterprise version and most likely is used for production services.

VMware guests hosted on the same ESX Server usually run in parallel. While this configuration provides great benefits for various types of applications, it might cause issues with backup and recovery. These issues are likely to occur when traditional backup techniques are used without careful planning and without consideration of the underlying equipment.

Figure 5-22 is an example of the simple ESX Server architecture. Two virtual servers (VMs) run on the ESX Server hardware. Their local drives are in fact virtual disks (that is, files) that lie on SAN-attached storage controlled by the ESX Server. The ESX Server virtualizes the virtual machines and makes them available in the network.

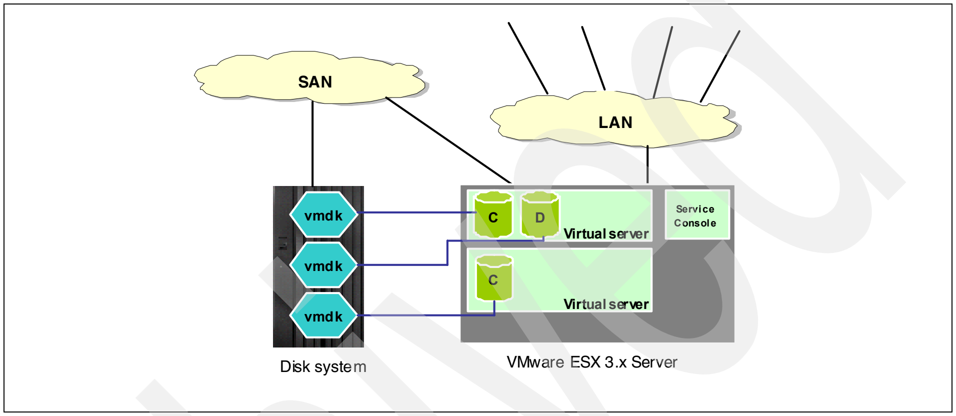

Figure 5-22 VMware ESX Server architecture

5.6.1 Current VMware environment

The Education Company has implemented VMware and benefits from it every day, seeing great reductions in hardware and power costs.

Figure 5-23 on page 264 shows the implementation of the virtualized servers. The VMware ESX Server consists of the individual guests with their respective hard disks stored on the DS8000 disk system.

The Education Company initially installed five VMware guests, which have now grown to 20 guests, and the daily backup routines are fast becoming a challenge. The workload on the VMware ESX Server during backup is substantial, and a second VMware ESX Server probably is required, even though the server is performing adequately during normal production. The Education Company is also worried that a restore of whole guests might take too long in case of recovery after a disaster.

> **Note:** Only two guests (VM1 and VM2) are shown in Figure 5-23 to simplify the illustration. To further simplify the backup and restore scenarios discussed in later sections, we have the Education Company running Microsoft Windows 2003 on their guests. The company uses no other operating systems.

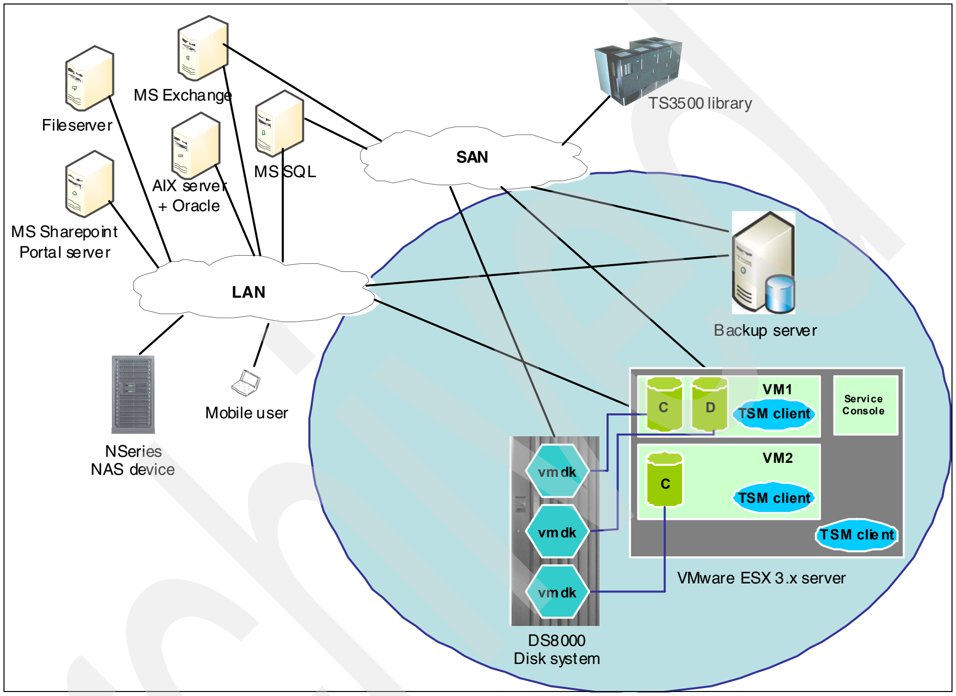

Figure 5-23 Current VMware environment

Current backup and restore strategy

At the moment, the Education Company uses the procedures discussed in the sections that follow.

Backup

The Education Company uses the following backup strategy:

- A backup client is installed on each of the 20 VMware guests. This client functions exactly as a backup client installed on a physical server. Regular incremental backups are performed, and backup data is sent over the LAN to the backup server.
- For faster recoverability of whole guests, a backup client is also installed on the VMware ESX Server itself. This client backs up the individual vmdk files

that constitute the guests' virtual disks. Before backup, the VMware ESX Server takes a snapshot of the vmdk file. The VMware ESX Server makes sure all Microsoft Windows disk buffers have been flushed, so the snapshot is consistent. A backup of the snapshot is performed and sent to the backup server across the LAN. Each backup is very large (several GBs) because all the guests' data is backed up.

Restore

The Education Company uses the following restore strategy:

- Restoring individual files is done from the file-level backups.
- The vmdk backup from the VMware ESX Server is used for restoring one or more guests and enables fast recovery.

The flow of backup and restore data can be seen in Figure 5-24. Notice that all data moves across the LAN and directly to and from the VMware ESX Server.

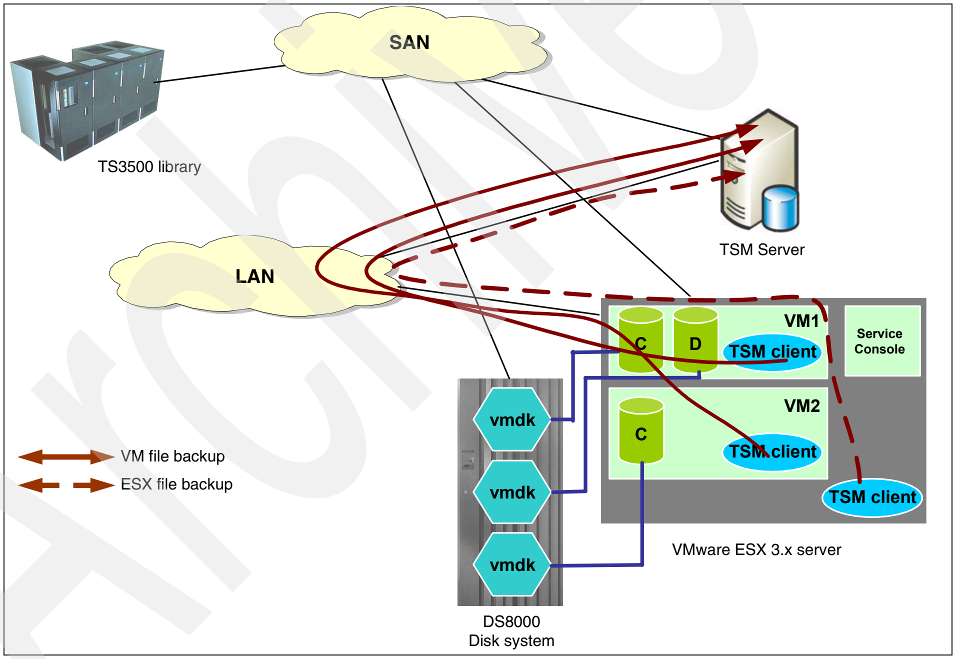

Figure 5-24 Flow of data

The current strategy has the following disadvantages:

- Running multiple Tivoli Storage Manager clients simultaneously on the same physical hardware (in different guests) uses up too much of the VMware ESX Server resources. It affects guests running production applications as well as guests running backups. When concurrent backups are running, the load on CPU, RAM, and disk I/O highly intensifies. Even single backups put an unacceptably high load on the VMware ESX Server.

- Configuring backup scheduling to avoid concurrent backups from the same VMware ESX Server is proving complex and hard to maintain and scale.
- Installing and customizing 20 Tivoli Storage Manager clients requires too much administrative effort, both initially, and whenever software upgrades are required. Of course, this problem affects nonvirtualized environments too, but the difference is more apparent in an automated VMware environment where builds are otherwise relatively fast and simple.
- LAN-free backup is not supported for VMware guests or for the VMware ESX Server, making all backup and restores go through the LAN. This configuration does not provide adequate performance for restores of vmdk-based backups.
- An incremental approach for the vmdk-based backups is lacking, so full backups of whole guests are performed each time.

The Education Company wants to implement a different backup and restore strategy. They can no longer afford depleting their VMware ESX Server resources during backups. In addition, recovery times for whole guests are not satisfactory.

5.6.2 Recommended VMware environment

The Education Company chooses to implement VMware Consolidated Backup (VCB) to handle backup and restore of their VMware guests. Along with native integration for VCB in Tivoli Storage Manager V5.5, this combination enables their previously discussed requirements to be met.

VCB is the preferred approach to backup in VMware for production environments. The fundamental idea behind VCB is to offload the backup workload from the VMware ESX Server to a dedicated, nonvirtualized system known as the *proxy node*. With proxying, the proxy node performs the backup operations on behalf of the guests, therefore offloading the VMware ESX Server. VCB is a component of the VMware infrastructure and, at the time of writing, is exclusively available for guests hosted on VMware ESX Servers.

The proxy node is a Microsoft Windows 2003 Server with the VCB Framework installed (the VCB Framework is a VMware software component). VCB Framework allows the proxy node to initiate snapshots and to withdraw them after use. To do so, VCB Framework communicates over TCP/IP to the VMware ESX Server.

A benefit of proxying is the ability to do LAN-free backups of guests' data. This ability is possible through the use of Windows 2003 on the proxy node because the Tivoli Storage Manager Storage Agent is supported on this OS.

Components

Figure 5-25 on page 268 illustrates the integration of VCB in the Education Company environment.

- A proxy node is installed with Microsoft Windows 2003, VCB Framework (for communication with the VMware ESX Server), an IBM Tivoli Storage Manager client V5.5, and an IBM Tivoli Storage Manager Storage Agent V5.5. The Tivoli Storage Manager client handles all data movement from the proxy node to the Tivoli Storage Manager server over the LAN. The storage agent handles all SAN traffic when LAN-free backups are running.
- The Tivoli Storage Manager server handles all backup data sent from the proxy node. File-level data is written to a disk-based storage pool, while SAN backup is written directly to tape in a TS3500 library.
- The VMware ESX Server hosts two virtual machines (or guests). A number of vmdk files are allocated to each guest to store data. The vmdk files reside in the DS8000 disk system and are zoned, so both the VMware ESX Server and the proxy can access the disks. No Tivoli Storage Manager clients are installed on the guests and the VMware ESX Server. All backup operations are handled on the SAN or LAN level through the proxy node.

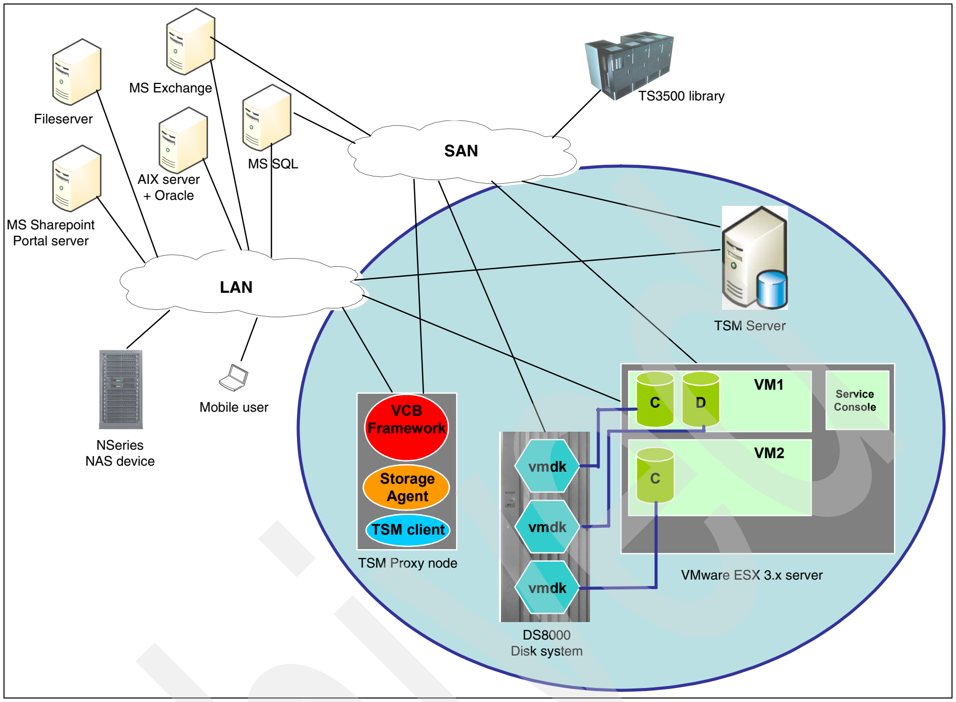

Figure 5-25 Integration of VCB backup with Tivoli Storage Manager

File-level backup and restore

In this section we take a look at the new file-level backup and restore options.

Backup

File-level backup allows the file systems of Microsoft Windows guests in the VMware ESX Server to be presented across the SAN to the proxy node, as shown in Figure 5-26 on page 269.

The VCB Framework software makes a snapshot on the VMware ESX Server of the file system data of the Microsoft Windows guest and virtually mounts it across the SAN into a folder on the proxy node's file system. In Figure 5-26 on page 269, the guests' C: drive is mounted to E:\VM1\letters\C\ on the proxy node. Although the folders look like local files, they are actually mounted as a virtual mount point on the proxy node. Data is not copied anywhere, only mounted, and therefore takes up no space on the proxy node.

The data is available to applications on the proxy node, just as it would be from ordinary mount points. The Tivoli Storage Manager client is able to access the data and do normal incremental backups to the Tivoli Storage Manager server across the LAN.

After backup, the file spaces created are associated with the actual guest's node name (not the proxy node name), and the Tivoli Storage Manager database therefore records and expires these files individually. Files and other objects appear as though they belong to the Tivoli Storage Manager node registered for that guest (not to the proxy). Thus from a Tivoli Storage Manager server perspective, each guest looks like it has been backed up from a locally installed client on the guest.

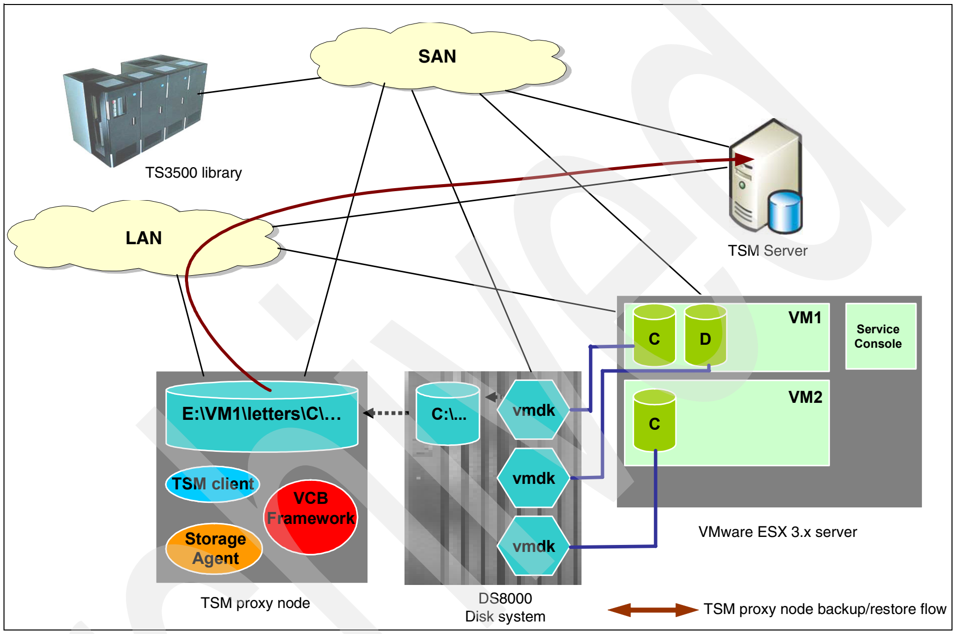

Figure 5-26 File-level backup and restore

Restore

You can use one of the following four methods when restoring data backed up to Tivoli Storage Manager. Each depends on both the destination for the restored files and which client node actually performs the restore.

- Restore files directly to the guest. This method requires that you install the Tivoli Storage Manager client on the guest to initiate the restore. Because VCB does not require the guest to have any Tivoli Storage Manager client code for backup operations, the client must be installed before the restore can run. When the client is installed, you can browse for and select files for restore as usual.
- Restore files to the proxy's local disk. If you do not want to install the Tivoli Storage Manager client on a guest, you can restore files to the proxy node's local disk, using the existing Tivoli Storage Manager client. This commonly used method of recovery works effectively for small amounts of data or

numbers of files. Once the files are on the proxy node, they can be accessed there or copied to the guest. A benefit of this method is that installing a Tivoli Storage Manager client on the guest itself is not required, thus saving disk space and administration costs.

- Restore to an alternative machine than the proxy or the guest. To restore to another machine, you must first grant access to the guest's backup data, after which the files can be restored to the alternative machine.
- Restore to the guest with the proxy node and a CIFS share. The proxy node can restore files directly onto the guest by mounting the guest's file system (with appropriate permissions) as a CIFS share. You might lose some of the NTFS security ACLs using this method, depending on the configuration of your file sharing system.

Full backup and restore of guests

This section describes how to use Tivoli Storage Manager to back up and restore a full Virtual Machine. When referring to backing up and restoring a full Virtual Machine, we are referring not to the individual files inside the guest but to something similar to a image backup of the whole guest system.

In this section, we solely address LAN-free backups. You can also perform a full Virtual Machine across the LAN, but the significant performance benefit of backups across the SAN is lost when you do so.

Backup

Full Virtual Machine backup is particularly applicable to disaster recovery environments because it is quite simple and fast to recover full Virtual Machines using this technique. LAN-free movement of full Virtual Machine backups is particularly useful for speedy recovery. In addition, it removes the backup load from the Tivoli Storage Manager server.

A significant difference between full Virtual Machine and file-level backup is that the full Virtual Machine backup requires and uses a large amount of temporary disk space on the proxy node. By comparison, a file-level backup uses no space locally on the proxy node because it simply mounts the same volumes used for the running guests on the proxy node and caches all changes on the VMware ESX Server. With full Virtual Machine, the data is physically copied to the proxy node, not just mounted on the proxy node.

The full Virtual Machine temporary disk space must be at least as large as the disk image being backed up. In other words, the full Virtual Machine disk space must be large enough to store all the disks from the largest guest within its temporary space. This space is occupied by the full Virtual Machine disk slices while the backup is in progress and is deleted after each backup. We recommend you use temporary space on a drive other than the proxy's system drive, so the stability of the proxy node can be maintained during operation. In fact, we recommend using a fast, reliable disk for the process because it is very I/O intensive.

The full Virtual Machine backup process copies the disk image from the VMware ESX Server to the proxy node's temporary space over the SAN and produces files of manageable size (2 GB per image slice by default, although you can adjust the size). This process is depicted in Figure 5-27 on page 272. These files, which include the VMware vmdk files, can then be backed up by the proxy node to the Tivoli Storage Manager server using a normal selective backup, as though the files were the proxy node's own files.

Incidentally, we use selective backup as opposed to incremental backup with full Virtual Machine images. We do so because when we use subfile backup, all the subfile chunks can be retained for the same amount of time. full Virtual Machine backup works well with Tivoli Storage Manager adaptive differencing (subfile backup) technology, which eliminates much of the backup overhead of taking full images at the client side, before they ever make it to the Tivoli Storage Manager server. This process facilitates very efficient backups both from a client-processing perspective and from the perspective of overall storage utilization on the Tivoli Storage Manager server. Only changed blocks are transferred to the Tivoli Storage Manager server during subfile backup.

Alternatively, files can be archived. Archiving the files groups the files together and makes it easier to locate them for subsequent retrieves. A drawback of archives is the inability to use adaptive differencing technology, thus having to transfer all files to the Tivoli Storage Manager server during every single archive.

We strongly recommend preproduction testing of VCB solutions in order to more accurately predict resource requirements for your future production environment.

Note: In theory, any supported guest operating system running under ESX, not just the Microsoft Windows guests, can be backed up using full Virtual Machine. We must point out that VMware Tools, which control the synchronization of the file system, is available only on Windows; therefore, other methods are required to achieve synchronization in non-Windows environments to ensure a consistent full Virtual Machine backup or archive.

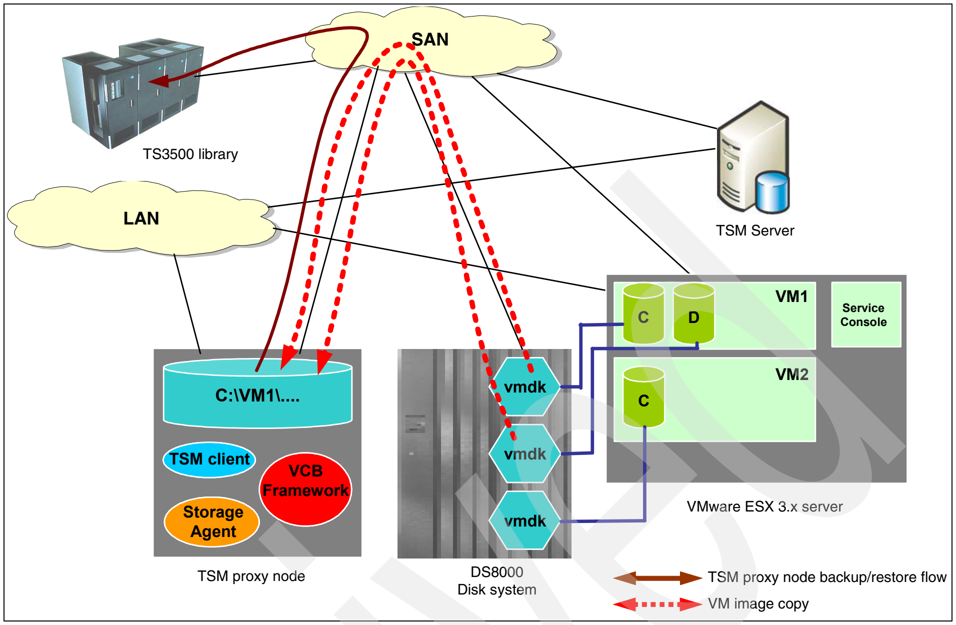

Figure 5-27 Full Virtual Machine backup and restore

Restore

This section discusses the restore of a full Virtual Machine process. As a first step, restore all the image files to a directory, for example, on the proxy node. Make sure you restore all the files from the same backup because mixing the files from different full Virtual Machine snapshots causes unexpected behavior. You can choose to restore the files based on the date they were backed up.

If you have archived your full Virtual Machine files instead of backing them up, simply retrieve all of them from the appropriate archive.

Once you have restored the files, you can then use the VMware Converter utility on a system of your choice to import the restored set of image files back onto a VMware ESX Server. You might choose to put the VMware Converter on the proxy so that you back up and restore from the same system. If you choose to do so, make sure you restore the image files to an location other than the backup mountpoint directory, so that you do not interfere with scheduled backup operations.

The VMware Converter can be downloaded from the following Web site:

http://www.vmware.com/products/converter/

Single file restore is not possible directly from the full Virtual Machine backup images. The entire image has to be restored if you need to extract single files.

> **Note:** It is possible to restore individual files using a VMware command-line utility (`mountvm`). The utility mounts a vmdk disk file as a virtual disk on top of an existing hard drive's file system. Once it is mounted, you can browse the VMware virtual disk and copy files from it as required.

Recommended VMware backup and restore strategy

The Education Company uses the new backup and restore functionality of VCB and Tivoli Storage Manager V5.5 and achieves the functionality they have been looking for.

The Education Company uses the following strategy:

- An incremental file-level VCB backup of all 20 guests is scheduled to run every day. That way the ESX Server is offloaded during backup, and production is not affected. The need for a second VMware ESX Server to share the load during backups no longer exists, and the investment in new hardware can be put off until more guests are added to the environment and the production load demands more resources. The company uses the proxy node when restoring small amounts of data or a small number of files, after which the files are copied to the guest in question. If more data must be restored, the company temporarily installs a Tivoli Storage Manager client on the guest and uses it for direct restores to that guest.
- Full Virtual Machine archives are run every Friday evening after business hours. The Education Company uses these archives for disaster recovery purposes. Because storage space in the IBM TS3500 library is of no concern, archives have been chosen as the preferred method of storing full Virtual Machine on the Tivoli Storage Manager server. They provide the greatest recovery speed, and the archived files are grouped together, making retrieves easier for inexperienced support personnel.

5.7 Summary

This chapter discussed some scenarios of using the new features of Tivoli Storage Manager V5.5. We discussed some aspects of the security implementation, an NDMP implementation, and a VMware backup implementation. We also discussed Microsoft SQL Server and Microsoft SharePoint.

6

Best practices, hints, and tips

In this chapter we discuss Tivoli Storage Manager software support and maintenance, performance tuning, hints, and tips to enable you to get the most out of your Tivoli Storage Manager server.

This chapter includes the following topics:

6.1 Software support and maintenance

This section deals with aspects of software management. To maintain a healthy Tivoli Storage Management environment, you must take care to ensure maintenance and patches are applied within a reasonable time frame.

Software management, in this case, deals with how you apply changes to your Tivoli Storage Manager environment.

You must consider the following types of changes when referring to a software management cycle:

- New products or agents are often introduced within the family of products, such as the IBM Tivoli Storage Manager HSM for Windows.
- A new version of an existing product is introduced that contains new features.
- A maintenance patch might be introduced to take advantage of new features that are present within a hardware device.
- Finally, interim fixes are released to correct an existing problem that was flushed out after the previous release of the product.

As stated previously in our discussion of the management cycle, change is probably inevitable with any software environment. You must also keep in mind that program support for older releases and versions of software is regularly being withdrawn. To avoid being in a position where a critical business environment is no longer supported, systems must be kept at current levels.

No matter how often new code levels are released, care must be taken within an environment to ensure that a large level of risk is not being introduced. As stated earlier, maintenance and patches must be applied within a reasonable time frame, but in some cases, you might be able to justify remaining at the current level within the environment. Some conditions require a complete upgrade to another version, conditions such as compatibility issues or options within the newly released code level not being fully accepted within standard practice.

It is beneficial to get into the habit of staying informed about new and upcoming releases and fixes for all areas of Tivoli Storage Manager. We offer the following resources to help you keep informed:

- IBM Software Support Web page:

  `http://www.ibm.com/software/support/index_A_Z.html`

  This location provides a complete list of all IBM-related software. You can find the Tivoli Storage Manager items listed alphabetically under the IBM Tivoli Storage Manager section. From there, you can select options to research current levels of all aspects of the Tivoli Storage Manager family.

- IBM FTP Index of Tivoli Storage Manager page:

  http://ftp.software.ibm.com/storage/tivoli-storage-management/

  Within this FTP location, which is the final destination for the software code in the IBM Software Support page previously listed, it is possible to directly obtain code without having to browse through all of the Web material. Readme files are located within each patch level, and they contain information about the fixes and requirements from an operating system (OS) and hardware point of view to enable you to make an educated decision of the version required.

- Automated e-mail notification of new products and releases, including information regarding versions being withdrawn from support:

  http://www.ibm.com/software/support/einfo.html

  Due to some compatibility issues with patch levels, you might be required to ensure that hardware within the environment is also at the latest firmware and driver levels. You can obtain various hardware and driver support at the following Web pages:

  http://www.ibm.com/servers/storage/support/allproducts/index.html

  http://ftp.software.ibm.com/storage/devdrvr/

6.2 Performance tuning

IBM Tivoli Storage Manager performance can be influenced by various tuning parameters. Diligence and expertise are required to tune these functions for good performance. IBM Tivoli Storage Manager can be quite complex because of the many OSs, network configurations, and storage devices that Tivoli Storage Manager supports. It is the purpose of this tuning section to provide hints and tips to ease the task of achieving desirable performance for Tivoli Storage Manager operations.

Tivoli Storage Manager functions in the client/server realm support many OSs, work across networks, and accept different communication protocols. Consequently, many other factors can also affect performance. The following factors can affect Tivoli Storage Manager performance quite significantly:

- Average client file size
- Percentage of files changed since last incremental backup
- Percentage of bytes changed since last incremental backup
- Client hardware (CPUs, RAM, disk drives, network adapters)
- Client OS
- Client activity (non-Tivoli Storage Manager workload)
- Server hardware (CPUs, RAM, disk drives, network adapters)

- Server storage pool devices (disk, tape, optical)
- Server OS
- Server activity (non-Tivoli Storage Manager workload)
- Network hardware and configuration
- Network utilization
- Network reliability
- Communications protocol
- Communications protocol tuning
- Final output repository type (disk, tape, optical)

This section discusses the following topics:

6.2.1 Measuring performance and discovering bottlenecks

Fine-tuning Tivoli Storage Manager is beneficial not just for backups but for restores as well. Some parameters might improve backup at the expense of restore and vice versa. Depending on the frequency and criticality of each backup and restore operation, you must determine the best balance within the environment.

Performance tuning and bottleneck discovery are complex tasks requiring thorough knowledge of Tivoli Storage Manager, the underlying OS, and the attached infrastructure components. Remember that a high-performance environment does not exist without any bottlenecks. To get on the right path with minimum constraints, you must discover and reduce the existing bottlenecks.

Network benchmarking

Network benchmarking within Tivoli Storage Manager is the process of measuring a given network’s capacity at a specific time during the day or evening.

At times, during a Tivoli Storage Manager deployment, you might discover that the product is not performing as well as expected. Perhaps you encounter slow backup and restore times. You must understand that all backup products are essentially an integration of bottlenecks. As a result, you must be able to isolate and address these weak links as you identify such issues.

You can follow a few procedures to identify some issues that might be present within the environment. The easiest method is to utilize basic file transfer utilities to determine a baseline of the server data transfer rates. Gather output for both pushing and pulling the data from a specific server, with the Tivoli Storage Manager server being the target. Make note of the file parameters.

Use one larger file as opposed to a directory source of smaller files equal in size. The best solution is to utilize a compression application to combine or compress (or both) a number of smaller files into one big file of the required size.

The same file used in the transfer earlier is then used to perform a selective backup from the same server from the Tivoli Storage Manager client. When dealing with just a single large file, you do not see any database processing overhead, but can collect comparable output for network throughput. You can then use this process in the same way for a restore using the replace parameter.

Tracing

If a noticeable difference is present between the push and the backup or the pull and the restore, you must investigate further. You can do so using the Tivoli Storage Manager tracing facilities that are present within the respective server and client agents. For further information regarding the use of the testflag and instrumentation options, refer to the problem determination guide for the client and server version in use. Trace routines require processing time and can slow down the response times of the server applications. Call a Tivoli Service Representative for assistance in diagnosing the output.

6.2.2 Infrastructure performance tuning

This section discusses considerations from a infrastructure design point of view.

Database and recovery logs

When implementing the Tivoli Storage Manager server, multiple physical disk volumes are beneficial for the database and recovery logs. Multiple volumes enable the server to spread its I/O over several volumes in parallel, which increases read and write performance. Subsequently, a greater quantity of smaller disks can provide better performance compared to a fewer number of larger disks of the same rotation speed.

To support the this configuration, it might be possible to stripe the database and log volumes across multiple disks. Doing so balances the I/O load so that it is spread not only over multiple database volumes but also over multiple disks.

Note: When spreading database and log volumes over multiple disks, ensure that the DB volumes do not share the same physical disks as the LOG volumes. When including mirroring for these volumes, again ensure that separate physical disks are also chosen for the mirrors. The remaining unused disk space on these volumes can be used to spread out the disk pools, again ensuring balanced I/O.

Another option that supports the previously described configuration is putting multiple database volumes on a striped volume set of disks to increase performance. By locating multiple database volumes on separate disk volumes, mirroring can increase not only server stability, but also server performance. The Tivoli Storage Manager server can access multiple database copies at the same time and prioritizes the volume copy with the shortest response time. Thus, if hardware does not support mirroring, the Tivoli Storage Manager server does very well, provided that I/O is properly balanced with the appropriate volume placement.

Tivoli Storage Manager access to the database for reading and writing is highly random in nature. Keep in mind that large read-ahead buffers defined for the volumes might hurt performance rather than help.

Storage pools

The options discussed in the previous section are also valid for disk storage pools, that is, performance is increased by using multiple small volumes for disk storage pools instead of using one large volume.

It is also preferable to send backups of large amounts of data, such as entire client databases, directly to tape instead of buffering in a disk storage pool. Using tape for this purpose allows the client to take advantage of the high streaming performance of most advanced tape drives when writing large files. In addition, it avoids extra data movement from the migration process, which is probably triggered immediately.

Backups of file servers and workstations must be directed to a disk storage pool to avoid tape mount delays. Also, a disk storage pool supports more multiple sessions than the available tape devices.

Versioning

Keeping multiple backup versions of files makes the Tivoli Storage Manager server database quite large. A larger server database cannot be accessed as quickly as a smaller one and requires more time for backup and inventory expiration operations. The planning process is important in order to make sure that only required versions of the data are processed.

Incremental backups are generally preferred over selective backups, because selective backups generate a new version every time, regardless of whether the specific file has changed.

Client configuration

By utilizing the powerful include and exclude options from either the client or server side of the configuration, you can specify at a detailed level the files to be backed up. Careful use of these options can reduce the daily amount of backup data significantly and thus improve both the client and server performance.

LAN-free backup and restore

Using LAN-free backups can greatly decrease LAN traffic and reduce the overhead on the Tivoli Storage Manager server, because the data is directly transferred to tape. This feature is important, especially for clients where backups, such as database backups, are routed directly to tape.

6.2.3 Tivoli Storage Manager server tuning

This section covers the most common tuning parameters for all Tivoli Storage Manager servers.

Server database and recovery logs

You can use the following parameters to tune the server database and recovery logs:

- BUFPOOLSIZE

  Cache storage is provided by the database buffer pool size, which allows frequently used database pages to remain cached in memory. This capability means faster access times and better performance because the pages are read and updated in memory instead of requiring I/O operations to external storage. The larger the database buffer pool, the more space for keeping pages in memory. However, because the buffer pool is implemented in real memory, you must balance its size against total memory requirements for the complete system and any other workloads.

  To measure the effectiveness of the database buffer pool, you must monitor the cache hit percentage. This number indicates how often a database page request can be satisfied from the cache as opposed to requiring disk access. An optimal setting for the database buffer pool is one where the cache hit percentage is greater than or equal to 98%. You can use the **`query db f=d`** command to check this value.

  Increasing the BUFPOOLSIZE parameter can improve the performance of many Tivoli Storage Manager server functions, such as multiclient backup, storage pool migration and backup, expiration processing, and move data processes. If the cache hit percentage is lower than 98%, increase the BUFPOOLSIZE value in the server options file. For most servers, we recommend starting with a value of 32768, which equals 8192 database pages. If you have enough memory, use increases in 1 MB increments. A cache hit percentage greater than 98% indicates that the proper BUFFPOOLSIZE has been reached. However, continuing to raise the BUFPOOLSIZE beyond this level might still be helpful. While increasing the BUFPOOLSIZE setting, take care not to cause paging in the virtual memory system. You must monitor system memory usage to check for any increased paging after changing the BUFPOOLSIZE.

Tip: Tivoli Storage Manager can use the parameter SELFTUNEbufpoolsize YES in the server options file to activate self-tuning. When you activate this option, the server checks the cache hit ratio each time inventory expiration is run. If the value is less than 98%, the option increases the BUFFPOOLSIZE by a predictable factor related to the actual cache hit ratio achieved.

- MIRRORWRITE

  Use this parameter to specify how mirrored volumes are accessed when the server writes pages to the recovery log or database during normal processing. A value of sequential specifies that the server writes to mirrored volumes one after another. Server performance decreases because a successful I/O must occur before the next write can occur. A value of parallel specifies that the server writes to all mirrored volumes in parallel. This mode allows all mirrors to obtain the new page at approximately the same time.

  Note: If a system outage occurs at the instant each mirror is partially complete in writing its page, a partial write to each mirror can result. See the DBPAGESHADOW option to prevent partial writes to each mirror.

  When you are using mirrored database volumes in sequential mode, a disk system with write caching ability can improve database performance. When writing in sequential mode, one database volume write must complete before the mirror is written to. A disk system with a write cache allows the mirror to be written to almost instantaneously, nearly doubling the amount of write activity possible. Many disk systems can do so although you might have to specifically enable them. Check your disk vendor documentation for details.

- DBPAGESHADOW

  Use this parameter to specify whether database page shadowing is enabled. If this parameter is enabled, Tivoli Storage Manager mirrors every write to a database page. This mirror is represented by a shadow file that contains these written mirrors. DBPAGESHADOW YES eliminates the requirement to set MIRRORWRITE SEQUENTIAL to prevent exposure to partial page writes. The default option for this parameter is NO.

Over time, database volumes can become fragmented. The efficiency of the database can be recovered and performance can be improved by reorganizing the database. By performing a database unload and reload process, the database is compressed and reorganized.

Transactions

You can use the following parameters to tune server transactions:

- TXNGROUPMAX

  Use the TXNGROUPMAX option to specify the number of files that are transferred as a group within transaction commit points between a client and a server. A group of client files is known as an *aggregate* or a *physical* file. Client backup, archive, restore, and retrieve performance might improve by using a larger value for this option. The default value for this parameter is 256 files.

  TXNGROUPMAX is used in conjunction with the client TXNBYTELIMIT option.

- MOVEBATCHSIZE and MOVESIZETHRESH

  You can use the MOVEBATCHSIZE and MOVESIZETHRESH parameters to help tune the performance of the server processes that involve the movement of data between storage media. These processes include storage pool backup and restore, migration, reclamation, and move data. The options specify the number of files and the quantity of data that are to be moved and grouped together in batch within the same server transaction.

  MOVEBATCHSIZE and MOVESIZETHRESH apply to the file aggregates or physical files as previously described in our discussion of the TXNGROUPMAX parameter. MOVEBATCHSIZE and MOVESIZETHRESH have default values of 1000 files and 2048 KB, respectively.

> **Note:** You can configure Tivoli Storage Manager to autoconfigure and monitor these parameters so that it can increase the range for better throughput.
>
> When you set the SELFTUNETXNSIZE parameter to YES in the server options file, you can set both the MOVEBATCHSIZE and MOVESIZETHRESH server options to their maximum values for optimal server throughput. Also, the server monitors and changes the value of the TXNGROUPMAX parameter to optimize client/server throughput.

With the TXNGROUPMAX and TXNBYTELIMIT parameters in use, more recovery log space might be required due to large transactions. Using these parameters might also mean that the Tivoli Storage Manager server requires more time to process log entries on restart.

Communication

You can use the following parameters to tune server communication:

- MAXSESSION

  You can use the MAXSESSION parameter to specify the maximum number of simultaneous client sessions that can connect to the Tivoli Storage Manager server. The default value is 25 client sessions, and the minimum value is 2. The maximum value is limited only by available virtual memory or communication resources. By limiting the number of clients, you can improve server performance, but you reduce the availability of Tivoli Storage Manager services to the clients.

  Because Tivoli Storage Manager clients use multithreading, you must consider the total number of concurrent nodes times the number of concurrent sessions that the server must handle. If you require 10 concurrent nodes to perform backup operations and each of the nodes has a multithreading capability (resource utilization) of 4, you must set the MAXSESSION to at least 40 (10 x 4).

- MAXSCHEDSESSIONS

  Use the MAXSCHEDSESSIONS parameter to specify the number of sessions that the server can use for processing scheduled operations. You specify the number as a percentage of the total number of server sessions available (the MAXSESSION parameter). You must set this number to less than 100; otherwise, it can lead to a situation where all the available sessions are being used for running client scheduled backups, and a critical restore cannot run because no spare session slots are available.

  Use the **`query status`** command to display the current MAXSCHEDSESSIONS value and change it to another value with the **`set maxschedsessions`** command. The default setting has value of 12.

- TCPWINDOWSIZE

  Use this option to specify the size of the TCP sliding window in kilobytes. The TCPWINDOWSIZE option overrides the OS's TCP send and receive spaces. In AIX, for instance, these parameters are tcp_sendspace and tcp_recvspace. The TCPWINDOWSIZE option specifies the size of the TCP sliding window for all clients and all but IBM Multiple Virtual Storage (MVS™) servers. A larger window size can improve communication performance but uses more memory. It enables multiple frames to be sent before an acknowledgment is obtained from the receiver. If long transmission delays are being observed, increasing the TCPWINDOWSIZE might improve throughput.

The size of the TCP/IP buffer is used when sending or receiving data. The window size used in a session is the smaller of the server and the client window sizes. Larger window sizes use additional memory but might improve performance.

We recommend you set the TCPWINDOWSIZE to 64 (except in Microsoft Windows, where our recommendation is 63). Clients now default to 63.

> **Tip:** On AIX, you must set the network option rfc1323 to on in order to have TCPWINDOW sizes larger than 64 KB.

- TCPBUFFSIZE

  You can use the TCPBUFFSIZE option to specify the size of the internal TCP communication buffer that is used to transfer data between the client node and the server. A large buffer can improve communication performance but requires more memory.

  We recommend the default setting of the TCPBUFFSIZE of 32768 KB on the client and server.

- TCPNODELAY

  We recommend you set the TCPNODELAY option to YES (the default) to allow data packets less than the Maximum Transmission Unit (MTU) size to be sent out immediately.

General Parameters

You can use the following general parameters to tune the Tivoli Storage Manager servers:

- EXPINTERVAL

  The Tivoli Storage Manager server runs automatic inventory expiration to remove client backup and archive copies, which are not required, from the server. You can use the EXPINTERVAL option to specify the interval in hours between automatic processing.

  Backup and archive copy groups specify the criteria that make copies of files eligible for deletion from a data storage pool. However, when a file becomes eligible for deletion, the file is only marked for deletion in the server database. It is not physically deleted from the database and storage pools until expiration processing occurs. If expiration processing does not occur periodically, storage pool space is not reclaimed from expired client files, and the Tivoli Storage Manager server requires an increase in disk and tape storage space.

Expiration processing is very CPU intensive and, if possible, must be run when other Tivoli Storage Manager processes are not occurring. To enable expiration processing to run as recommended, either schedule expiration once a day, or use the **`setopt expinterval 0`** command and manually start the process with the **`expire inventory`** command at the server.

> **Note:** If you set SELFTUNEbufpoolsize to YES, **`expire inventory`** processing also triggers the buffer pool performance monitoring algorithm to check whether a performance improvement can be gained (based on previous server execution and available memory).

- THROUGHPUT thresholds

  Use THROUGHPUTTIMETHRESHOLD and THROUGHPUTDATATHRESHOLD to control how Tivoli Storage Manager handles slow performing client sessions. Slow performing sessions might pin your recovery log and result in slower overall server performance. These options enable you to specify thresholds that must be met by client sessions in order to be considered productive and avoid cancellation.

  Specify THROUGHPUTTIMETHRESHOLD in minutes; it is added to the media wait time to determine the total client session time. It specifies the time threshold for a session after which it might be canceled for low throughput.

  Specify THROUGHPUTDATATHRESHOLD in kilobytes per second, and specify the throughput that a client must reach to prevent being canceled after the time threshold is reached.

  If the server detects that the client data transfer rate is below the specified threshold, the session is considered to be eligible for termination after the time threshold is reached.

- COLLOCATION

  The speed at which you can restore the files, among other parameters, also depends on how many tape drives are available on the server and whether you are using collocation to keep clients nodes assigned to as few volumes as possible.

With tape capacities becoming larger and larger, collocation by groups of nodes is now supported. You can define groups of nodes, and the server can then collocate data based on these groups. Collocation by group can yield the following benefits:

- Reduce unused tape capacity by allowing more collocated data on individual tapes
- Minimize mounts of target volumes
- Minimize database scanning and reduce tape passes for sequential-to-sequential transfer

You can have the server collocate client data when the data is initially stored in server storage. Another option is if a storage hierarchy is present, you can have the data collocated when the server migrates the data from the initial storage pool to the next storage pool in the storage hierarchy.

For newly defined storage pools, the default storage collocation setting is now group.

6.2.4 Tivoli Storage Manager client tuning

This section takes a look at tuning parameters within the Tivoli Storage Manager client to obtain maximum performance when using Tivoli Storage Manager. You can change or add these parameters to the client options files or within a client options set and force them down from the Tivoli Storage Manager server.

Transactions

You can use the following parameters to tune client transactions:

- TXNBYTELIMIT

  Use this parameter to specify the batch size, in kilobytes, for Tivoli Storage Manager server transactions. TXNBYTELIMIT is used together with TXNGROUPMAX. The TXNBYTELIMIT option reduces the number of server transactions but increases the amount of data within any single transaction. Therefore, the amount of overhead during backup, restore, archive, and retrieve caused by database commits is reduced.

  A *transaction* is the unit of work exchanged between the client and the server. Because the client program can transfer more than one file or directory between the client and server before it commits the data to server storage, a transaction can contain more than one object, which is called a *transaction group*. The TXNBYTELIMIT option enables you to control the amount of data sent between the client and server before the server commits the data and changes to the server database, thus changing the speed with which the client performs work.

The amount of data sent within a single transaction applies when files are batched together during backup or when receiving files from the server during a restore procedure. The administrator can limit the number of files or directories contained within a group transaction using the TXNGROUPMAX option. The actual size of the transaction can be less than your limit. Once this number is reached, the client commits the transaction to the server even if the transaction byte is not reached.

If you increase this parameter, consider the following issues:

– More data per transaction increases server recovery log requirements. Check the log pool to ensure enough space is available and remember that a larger log might increase server startup time.

– More data per transaction might result in more data being retransmitted if a retry occurs. During a retry, all objects belonging to the current transaction must be retransmitted, which decreases performance.

– The benefits of changing this parameter are subject to configuration and workload characteristics. In particular, this parameter benefits tape storage pool backups more than disk pool backups, especially if the workload consists of many small files.

When setting the size of transactions, consider setting a smaller size if you are experiencing many resends due to files changing during backup when using static, shared static, or shared dynamic. This consideration applies to both static and shared because when a client detects that a file has changed during backup and decides not to send it, the client still has to resend the other files in the transaction.

- RESOURCEUTILIZATION

  Use the RESOURCEUTILIZATION client option to regulate the level of resources (the number of concurrent sessions) the Tivoli Storage Manager client and the server can use during processing.

  Tivoli Storage Manager clients are automatically multisession capable, and this function is transparent to the user. However, the RESOURCEUTILIZATION parameter allows the user to customize the function.

  This option increases or decreases the ability of the Tivoli Storage Manager client to create multiple sessions. The value of RESOURCEUTILIZATION does not represent the number of sessions created by the client; it specifies the level of resources the Tivoli Storage Manager server and client can use during backup or archive processing. During backup or archive, the higher the value, the more sessions the client can start if necessary.

  The parameter range is from 1 to 10. When the option is not set, the default is creating two sessions to the server. One session is opened for querying the server, and the other for transferring file data.

If you set the RESOURCEUTILIZATION option to enable multiple client/server sessions and are backing up direct to tape, you must also update the server MAXNUMMP option for the client node maximum mount points allowed.

The following factors can affect the throughput of multiple sessions:

– The server's ability to handle multiple client sessions (sufficient memory, multiple storage volumes, and CPU cycles to increase backup throughput?).

– The client's ability to drive multiple sessions (sufficient CPU, memory, and so on).

– The configuration of the client storage subsystem. File systems that are striped across multiple disks, using either software striping or RAID-5, can better handle an increase in random read requests than a single drive file system. A file system on a single drive might not see a performance improvement if it attempts to handle many random concurrent read requests.

– Sufficient bandwidth in the network to support the increased traffic.

The following potentially undesirable actions might occur when you run multiple client sessions:

– The client can produce multiple accounting records.

– The server might not start enough concurrent sessions. To avoid this result, review and possibly change the setting for the server MAXSESSIONS parameter.

– A query node command might not summarize client activity.

Communication

You can use the following parameters to tune client communication:

- DISKBUFFSIZE

  Use the DISKBUFFSIZE option to specify the maximum disk I/O buffer size (in kilobytes) that the client might use when reading files. Optimal backup, archive, or HSM migration client performance might be achieved if the value for this option is equal to or smaller than the amount of file read ahead provided by the client file system. A larger buffer requires more memory and might not improve performance. This option replaces the formerly available LARGECOMMBUFFERS option. We recommend keeping the default value unless otherwise advised by IBM support.

- TCPBUFFSIZE

  You can use the TCPBUFFSIZE option to specify the size of the internal TCP communication buffer used to transfer data between the client node and the server. A large buffer can improve communication performance but requires more memory.

- TCPWINDOWSIZE

  You specify the size of the TCP/IP sliding window in kilobytes with the TCPWINDOWSIZE. This option overrides the OS's TCP send and receive spaces. It specifies the size of the TCP sliding window for all clients and all MVS servers. A large window size can improve communication performance but uses more memory. TCPWINDOWSIZE enables multiple frames to be sent before an acknowledgment is obtained from the receiver. If you observe long transmission delays, increasing the TCPWINDOWSIZE might improve the throughput.

- TCPNODELAY

  When you set the TCPNODELAY option to YES, it allows data packets less than the MTU size to be sent out immediately. We recommend setting the value to YES.

General parameters

Use the following general parameters to tune the client:

- COMPRESSION

  The COMPRESSION option compresses files before they are sent to the Tivoli Storage Manager server. Compressing files reduces the amount of data storage required to store backup versions and archive copies of your files. It can, however, affect Tivoli Storage Manager throughput because of the CPU cycles required for the client to compress the file.

  Client data compression saves storage space, network capacity, and server cycles. If compression is on, backups can also be significantly slower, depending on client processor speed. Compression is beneficial for a client with a fast processor on a slow network but probably not for a slow client processor on a fast network.

  By default, files continue compressing even if the operation is causing the file size to increase. This result can be the case if you back up data that is already compressed. To prevent continued compression if the file size grows and to send the file again without compression, set the COMPRESSALWAYS option to NO. When this option is set to NO, the current transaction is rolled back, and the objects for the transactions must be retransmitted. You might want to edit your include/exclude list to exclude precompressed files from compression using the exclude.compression option.

The client COMPRESSION option is only in effect if you specifies that your client node might choose its own compression. When you register a node using the **`register node`** command, a COMPRESSION parameter allows the server to enforce compression or noncompression or allows the client to choose the setting. The default value is that the client chooses the setting within the client options file.

The following general rules apply to using Tivoli Storage Manager compression:

– If you are backing up directly to tape, and the tape drive supports its own compaction, use the tape drive compaction, not Tivoli Storage Manager compression. If you have already compressed the data at the client level, tape drive compaction is unable to reduce the data further.

– Do not use Tivoli Storage Manager compression if a client currently has built-in file compression support, such as the Tivoli Storage Manager client for NetWare. Tivoli Storage Manager compression on these clients does not yield additional reduction in the amount of data backed up to the server.

– If you are using compressed file systems, such as NTFS compression on a Microsoft Windows system, the Tivoli Storage Manager client reads the uncompressed data stream regardless of the compression state of that file in the file system. Thus, the file is uncompressed during the read and, if it is compressable, it is compressed again by the client during the backup. The decompression/compression overhead decreases backup performance. When the same file is restored, the client decompresses the file data received from the server and writes the uncompressed data to the file system, where it gets recompressed by the file system driver. Because of this required extra processing, we do not recommend using Tivoli Storage Manager compression on NTFS compressed file systems.

- COMPRESSALWAYS

  The COMPRESSALWAYS option controls what occurs when the compression actually increases the file size. You can continue compressing or send the object again. Use this parameter in conjunction with the COMPRESSION option.

- DIRMC

  This option had a major impact on restoring performance in earlier versions of Tivoli Storage Manager, where directories were restored first, followed by files. It was, therefore, an advantage to cache directory information so that it was permanently on disk. Best practice design was to create small storage pools specifically for directories. The DIRMC directive in the client options file was used to bind directories to a management class pointing to this storage pool.

Restore processing has changed since then, and directories are created with default attributes. The correct attributes and access control list (ACL) information are applied to the data as it is read from the backup media. Therefore the original reason to cache directories on disk no longer applies.

Nevertheless, the DIRMC option is still useful. If you do not specify this option to associate a management class with directories, the client uses the management class in the active policy set of your policy domain with the longest retention period, which could point to a storage pool on tape. This result might mean unwanted mount requests occur during a backup. Therefore, in the design described in this book (see Chapter 5, “Scenarios and best practices” on page 207), we have used a separate storage pool and management class for directories, and we recommend using the DIRMC option.

- QUIET

  The QUIET option keeps messages from being written to the window during backups. By default, Tivoli Storage Manager displays information about each file it backs up. To prevent displaying this information, use the QUIET option. Messages and summary information are still written to log files.

  The following benefits of using the QUIET option are the following two:

  – For tape backup, the first transaction group of data is always resent. To avoid this result, use the QUIET option to reduce retransmissions at the client.

  – If you are using the client scheduler to schedule backups, using the QUIET option dramatically reduces disk I/O overhead to the schedule log and improves Tivoli Storage Manager throughput.

6.2.5 System design for performance

The following sections provide guidance on designing solutions for performance.

PCI buses

Most IBM System p® machines have multiple PCI buses. For example, when you are backing up to tape over a LAN, it is not desirable to have the network adapter and the tape (SCSI, Fibre Channel) adaptor on the same bus. Spreading the PCI adapters across different buses based on predicted usage provides the most efficient layout.

Tape buses (SCSI, Fibre Channel)

Placing too many tape devices on a bus can cause contention issues. For HVD SCSI tape devices, place no more than two devices on a single bus.

For Fibre Channel tape devices, consider the throughput of the tape device and the HBA. For example, if you are using third-generation LTO drives with a rated performance of 80 MBps and your backup data streams are able to push the drives close to this limit, you might choose to put one or two drives per 2 Gbps HBA. Do not forget to consider the effect of compression.

It is often difficult to achieve the rated speed of high-speed adapters such as Gigabit Ethernet and Fibre Channel, no matter what the application. Do not expect to achieve the full rated speed with these cards.

Disk topology

It is desirable to have a single Tivoli Storage Manager volume per disk, be it database, log, or storage pool. This configuration provides the least contention, allowing each volume to act independently. This configuration is often difficult as disk sizes get larger and the use of RAID is insisted upon. It takes discipline to dedicate a 72 GB disk to a small 6 GB log volume, for example. Using virtualization technologies, such as the IBM TotalStorage SAN Volume Controller, can assist here because virtual disks can be made of any size.

In the case of RAID, you have little choice but to have multiple Tivoli Storage Manager volumes per logical disk.This choice is a built-in downside to RAID; instead of each disk acting independently, they must all work in unison, limiting performance.

When using JBOD, placing two storage pool volumes on a single disk does not result in a degradation and might provide a performance improvement. The performance improvement is because Tivoli Storage Manager writes synchronously to storage pool volumes, and providing the disk with other tasks to perform increases aggregate performance.

System memory

The target use of memory in a Tivoli Storage Manager server is the database buffer pool. A large buffer pool can improve database performance because of the reduced disk I/O required. You must provide enough real memory for an optimal buffer pool. The buffer pool is adjusted using the BUFPOOLSIZE option described earlier. Set the BUFFPOOLSIZE option to somewhere between one eighth and one half of the real memory.

Network

If possible, use dedicated networks for backup (LAN or SAN) so that other network traffic does not share the same wire as the backup data.

Gigabit technology supports jumbo frames (9000 bytes), which you must use whenever Gigabit Ethernet hardware is implemented. Unfortunately, not every Gigabit Ethernet hardware component supports jumbo frames. Using this feature is only practical when all the elements (server, switch, and client) support it.

Tape devices

Calculating your daily throughput for backup and archive data and the amount of data sent directly to tape must result in a minimum number of tape devices that must be in place to sustain your data flow. Remember to size the number of tape devices to serve the data flow during the peak backup window and not on a daily average basis.

If you have to schedule certain server operations, such as migration, reclamation, or storage pool backups, during the backup window, make sure that enough tape drives are available to support these additional operations. Even if you typically schedule these server operations outside the normal backup window, you must allow a number of tape drives serving unscheduled restore or retrieve operations.

6.2.6 Special performance options

The following sections focus on additional performance options.

LAN-free tuning

When using LAN-free backup and restore, you must set TCPNODELAY to YES (the default) on the storage agent and server.

LTO/DLT tape tuning

Due to the high latency of the **`write filemarks`** command, LTO and DLT tape devices benefit from making Tivoli Storage Manager transactions as large as possible. Setting the server TXNGROUPMAX to greater than 512 files and the client TXNBYTELIMIT to the maximum of 2097152 bytes can reduce the number of transactions, which enhances performance. It is important to understand that these settings can affect the server's recovery log utilization and increase the amount of time between transaction commits. In Tivoli Storage Manager V5.2 and later, the maximum setting for the TXNGROUPMAX value is increased from 256 to 65000. You must follow the double-up and test rule for these parameters.

File system volumes versus raw logical volumes

As a rule of thumb, we generally recommend you put Tivoli Storage Manager server volumes on *raw logical volumes* (*RLVs)* in a UNIX environment. In addition to the read/write performance, RLVs do not fragment, and creation is very fast. Support for RLVs is available on all Tivoli Storage Manager UNIX platforms, with the exception of Linux. RLVs are not supported for sequential-access storage pools (device class of type file).

Tests on AIX have shown little difference in performance between using Tivoli Storage Manager storage pool volumes on JFS2 file systems using direct I/O and RLVs. The same is true for tests on Solaris systems using VxFS file systems with the QuickIO option enabled.

> **Note:** On AIX, be aware that only storage pool volumes, not database or log volumes, attempt to utilize the AIXDIRECTIO option.

If you decide to place your server volumes on a file system, be aware of fragmentation in the file system, especially if you create a lot of volumes at the same time. Also, creating such volumes can take some time because the whole file is written during initialization.

AIX virtual memory system tuning

AIX is very aggressive in caching file system data. In many cases this feature is highly desirable. In the case of Tivoli Storage Manager, file system caching provides no benefit because the data is not likely to be in cache when it is required (for migration or restore). This lack of benefit is because the large volume of data Tivoli Storage Manager processes is typically much greater than available memory.

AIX often pages out application memory (including Tivoli Storage Manager) in favor of file system cache data. Paging out application memory is undesirable because it often causes the Tivoli Storage Manager database buffer pool to be paged out to disk, negating much of the benefit of the buffer pool. This result can cause misleading database cache hit statistics because Tivoli Storage Manager believes it has that data in memory, when in fact it has much of this data on disk in paging space. The symptoms for this behavior are high paging rates, which you can check with the **`vmstat`** command. This situation can be avoided by tuning AIX so that it does not favor the file system cache over application memory. You can tune AIX by using the **`vmtune`** command provided in the bos.adt.samples AIX file set.

By default, AIX uses up to 80% of memory as a file system cache. For machines dedicated to Tivoli Storage Manager, this amount must be lowered. A good starting point is 50%. You can lower the amount using the parameter maxperm. Run either the **vmtune** or the **vmo** command, depending on your version of AIX.

For a detailed explanation of the **vmo** tunables, see VMM page replacement tuning at:

http://publib.boulder.ibm.com/infocenter/pseries/v5r3/index.jsp?topic=/com.ibm.aix.prftungd/doc/prftungd/vmm_page_replace_tuning.htm

If, after a period of time, paging still occurs, you can further lower the amount of memory used as a file system cache. This behavior can also be avoided by using RLVs for disk storage pools. AIX does not attempt to cache data residing on a RLV and is then less likely to have to page out application memory. RLVs are also more efficient (in terms of CPU utilization) because AIX does not spend time attempting to cache this data. If a Tivoli Storage Manager server has high CPU utilization, better performance might be obtained by using RLVs.

System paging often is interpreted as a lack of memory. With Tivoli Storage Manager, this misinterpretation rarely does occurs, and paging can (and must) be eliminated by using the previously discussed tuning method. AIX clients can also benefit from altering the maxperm parameter. During backup, files are read into cache that are not likely to be read again.

As before, this behavior can cause application memory to be paged out, causing the main application of this machine to suffer during backup. Lowering maxperm can help. Beware of the effects of tuning maxperm on the client application because the performance of the main application might depend on file system cache in some instances (like a file server). In this case it might be desirable to simply alter maxperm before and after a Tivoli Storage Manager backup, or not at all.

Another tunable parameter is maxpgahead, which controls the AIX read-ahead algorithms. Read-ahead is desirable on the client during backup and on the Tivoli Storage Manager server during storage pool migration from disk.

We recommend setting maxpgahead to the maximum of 256. When tuning maxpgahead, you must also adjust the maxfree parameter to allow room for this data. The maxfree parameter must equal the minfree + maxpgahead (or greater).

For more information about the commands, see the AIX documentation or man pages for **vmtune**, **vmo**, and **ioo**.

Using NTFS partitions for the server

In a Microsoft Windows environment, place the Tivoli Storage Manager server recovery log, database, and disk storage pool volumes on partitions formatted using NTFS. NTFS has the following advantages:

- It offers support for larger disk partitions than file allocation tables (FAT).
- It has better data recovery.
- It has better file security.
- Formatting Tivoli Storage Manager storage pool volumes on NTFS partitions is much faster than on FAT partitions.

> **Attention:** NTFS file compression must not be used on disk volumes that are used by the Tivoli Storage Manager server due to the risk of performance degradation.

Journal-based incremental backup

Using journal-based incremental backup can provide substantially faster incremental backups of file systems with a large number of files and small daily change rates. You configure the journal options in the tsmjbbd.ini file in the client install directory. Refer to this file for supported settings and configurable parameters.

We recommend you leave the journal database size set to unlimited. Limiting the journal database size can potentially cause the journal to become invalid, and then Tivoli Storage Manager resorts to a full incremental backup. Of course, the journal size depends on how many distinct files and directories are updated or deleted between backups, and how much activity the file system sees in general. Very active file systems require more space for the journal database on a disk with at least 1 GB of free space. If free disk space is limited, consider performing backups more frequently than once per day.

6.2.7 Tuning summary

When all is said and done regarding performance, one thing must be clear: This job is never complete. As your client workload changes, your network configuration evolves, the amount of managed data increases, databases and their specialized client requirements change, and hardware gets faster. All of these factors combine to move bottlenecks around your environment in an endless cycle.

6.3 Housekeeping and reporting

This section describes general housekeeping and reporting hints and tips, including the following topics:

- 6.3.1, "Daily housekeeping" on page 298
- 6.3.2, "Tivoli Storage Manager reporting" on page 306

6.3.1 Daily housekeeping

This section describes the automation mechanisms Tivoli Storage Manager provides to initiate certain actions such as backups and housekeeping activities. We also discuss different means of communicating schedule information between the Tivoli Storage Manager server and clients.

Introduction

Tivoli Storage Manager includes a central scheduling component that allows the automatic initiation of administrative and client operations at predefined times. As an administrator, you are responsible for creating and maintaining the schedules in each policy domain.

Tivoli Storage Manager scheduling is divided into two categories, administrative scheduling and client scheduling. The two categories differ in the following three key areas:

Execution location	An administrative schedule performs an action on the Tivoli Storage Manager server, while the client schedule can only execute on a Tivoli Storage Manager client.
Domain privilege	Only an administrator with system privileges can manage an administrative schedule, while an administrator with policy privileges in the client's domain can manage the client schedule. This granularity can be useful when distributing management control across a large enterprise.
Commands	An administrative schedule can only initiate an internal Tivoli Storage Manager command, while a client schedule can initiate an internal client action such as an incremental backup, or run an external command, such as a shell script or executable.

For both types of schedules, the following four key pieces of information must be available:

- A command or action to be executed
- When the command or action executes

- The period, or window, in which the command or action must start
- How often the command or action must be repeated

The command or action that you run might be an incremental backup (client schedule) or a storage pool migration (administrative schedule) that you must run every day at a particular time. You also have to estimate how long the command may run so that you can synchronize your schedules and balance the load on the server.

For client schedules, you can specify whether the client must poll the server regularly to receive information about scheduled actions (client polling), or whether it must wait to be contacted by the server to start a scheduled action. This option is called the *scheduling mode* on the client, and you set it in the client options file.

Another type of client schedule, *Clientaction*, is for actions you want to run only once as opposed to recurring scheduled actions. The Tivoli Storage Manager server clock determines all schedule start times, regardless of the time zones where the clients are located.

Different *styles* of schedules are introduced in Tivoli Storage Manager V5.3. Both client and administrative schedules can now be either Classic or Enhanced. The styles refer to the way in which schedules are started and repeated. Classic refers to the original style of setting the start time and repetition of the schedule using a limited number of options. Enhanced refers to the new style, which provides more granular options for setting the repetition. You can now configure a schedule to occur, for example, on the last day of each month.

Figure 6-1 on page 300 shows a series of operations that can occur in a typical Tivoli Storage Management environment on a daily basis and the sequence of these operations. The circle represents a clock, and the sectors in the circle indicate the hours of the day. The daily schedule includes a period when clients perform their backups, from 10 p.m. until 4 a.m. the next morning. After the clients are finished, the server performs housekeeping.

In Figure 6-1 on page 300, the server makes copies of the disk storage pools for off-siting. The server backs up its database, deletes the volume history, saves the device configuration, and creates a list of tapes for vault processing. The server then migrates the data from the disk storage pools to on-site tape pools and reclaims blank space from the tape pools. Finally, the expiration process runs before a new round of client backups begin for that night.

Many factors influence the actual start time and duration of various operations, including the client backup window, storage pool sizes, amount of data, and so on. Nevertheless, you must carefully consider the timing and sequencing; if not, jobs can overlap and not complete properly, or jobs can tie up server resources unnecessarily.

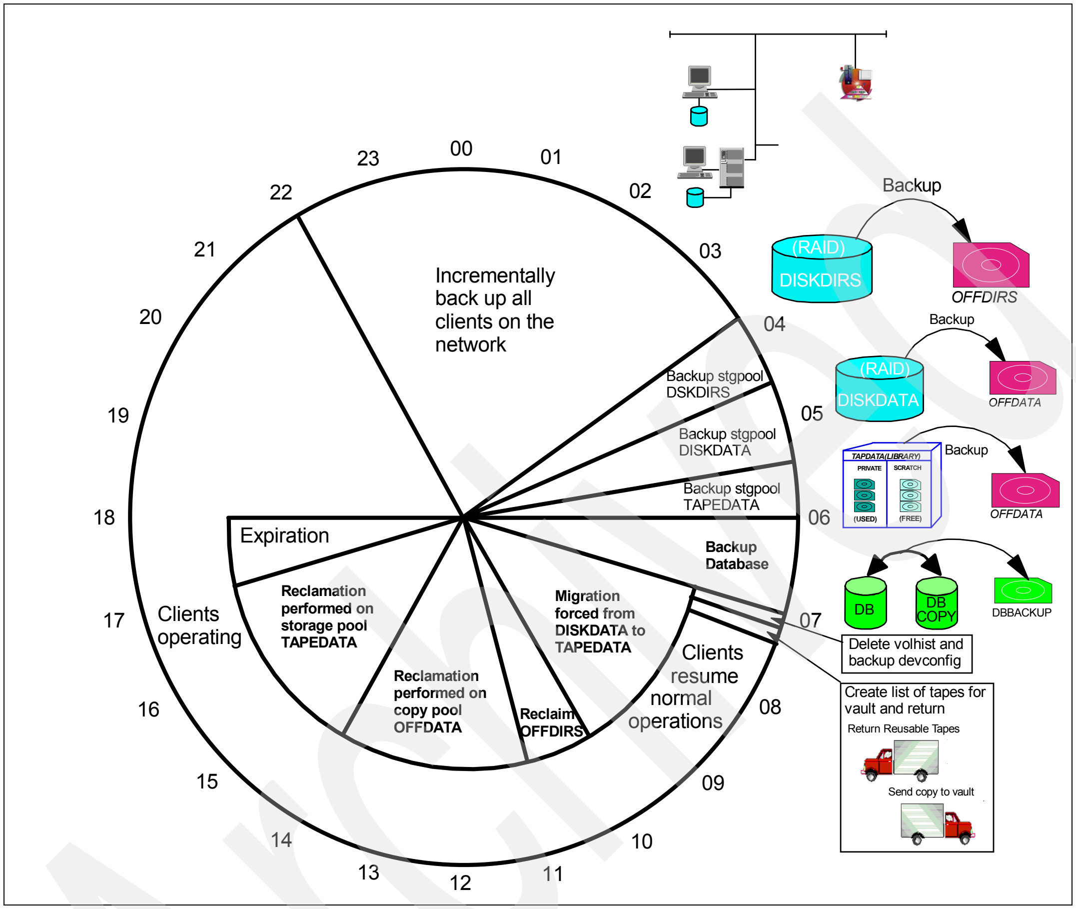

Figure 6-1 Client schedules

Administrative schedules

An administrative schedule is a directive to trigger an action within the Tivoli Storage Manager server. It consists of a server command and extra parameters describing when the action must take place. Because each administrative schedule can run only one server command, the command itself might be a **run** command, which runs an internally defined server script containing other internal server commands.

Scripting administrative commands can greatly assist timing and sequencing events; the scripting engine has directives to serialize commands and wait for them to complete before continuing processing.

You must define any actions that you perform on a regular basis to manage the Tivoli Storage Management environment as administrative schedules. Automating these operations to occur in a quiet period, such as overnight, enables you to ensure that server resources are available when clients require them.

Client schedules

A client schedule is a directive to trigger an action on one or more Tivoli Storage Manager client nodes. It is different from an administrative schedule in that it specifies an action to be performed on the Tivoli Storage Manager client. The client scheduling system consists of a server portion and a client portion. The server part is integrated into the Tivoli Storage Manager process and is responsible for maintaining the schedule parameters and tracking which nodes are associated with each schedule. The client scheduler is a separate process on the Tivoli Storage Manager client that communicates schedule information between the server and client. A client must be running its scheduler process to execute scheduled operations; otherwise, the operation is missed and is logged as such in the Tivoli Storage Manager server activity log.

The server event and activity logs record the success or failure of each scheduled operation. The administrator can query the logs to obtain the status of the schedules. The client also keeps a local log of scheduled operations.

The definition of the client schedule includes the actions to be performed. The action can be a single native Tivoli Storage Manager client command, such as a backup, a restore, an archive, or a retrieve command. The action also can be executing a *macro*, which is a collection of native commands, or an external command script. An *external command script* is simply any script that you can execute within the client OS, such as a Microsoft Windows batch command file, a UNIX shell script, or Perl script. The command script itself can contain Tivoli Storage Manager client commands plus logging, setup, error detection, post-backup procedures, and so on. Or you can use the script to schedule functions totally unrelated to Tivoli Storage Manager functions. The scheduler is therefore a flexible general purpose scheduler.

You define a schedule within the policy domain, so that only the client nodes belonging to that domain are eligible to execute this schedule. After defining the schedule, the administrator then specifies which client nodes (from those in the domain) execute the schedule. This action is called *associating* clients with a schedule. The administrator can choose to associate all of the nodes in the domain, or just a subset, according to requirements. The associations can be changed at any time.

Two scheduling modes are available for establishing communication between the client and server to start a scheduled event, which is shown in Figure 6-2. The selection of which mode to use is set in the client options file and is also dependent on the basic communication protocol used between the client and server. If the communication method is anything other than TCP/IP, then only the client polling method might be used. If TCP/IP is used, the client might select either method. Optionally, the server can override the client's preference if required.

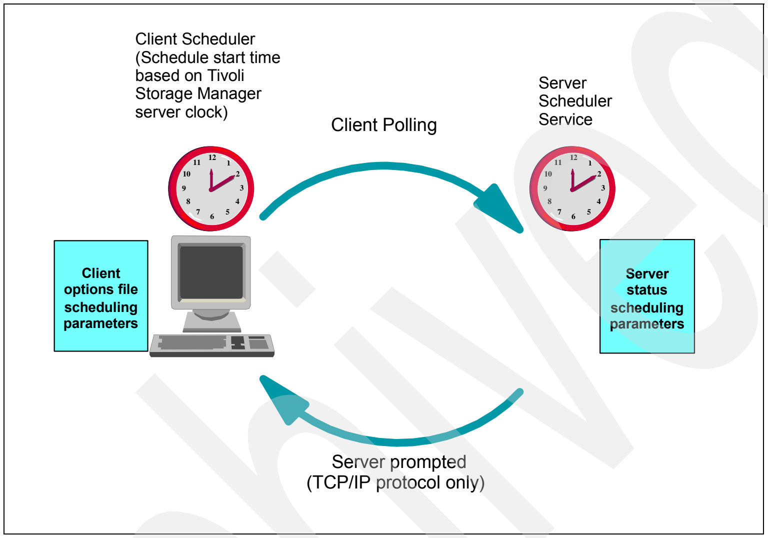

Figure 6-2 Client schedule types

Client polling

If a client has selected the client polling scheduling mode, the client contacts the server to discover if a schedule is defined for it and when it must be run. The client continues to contact the server at regular intervals set in the client's options file so that it can respond dynamically to any changes you have made to its schedules. You can also override this interval by setting a server parameter.

When the client contacts the server and a schedule is to be executed, the client receives a start time from the server. The client then counts down to the start time, and when the start time arrives, the client performs the actions defined by the schedule. If no action was scheduled to run between the time of contact and the next contact time, the client simply waits.

For example, a daily schedule might back up all of the client's file systems incrementally. Figure 6-3 on page 303 shows a client polling schedule configuration in which the client regularly queries the Tivoli Storage Manager server about the next operation to run. In this case, the Tivoli Storage Manager server informs the client when the schedule must run, and if it is not the scheduled time, the client waits. If the next poll is due to occur before the schedule is due to run, the client polls the server again at the time specified by the interval (QUERYSCHedperiod client parameter). If nothing has changed since the last poll, the server simply responds with the same information.

The main advantage of the client polling mode is that the server automatically can assign random start times for each client's schedule execution within a proportion of the schedule's start window. Randomized start times are useful if many clients have to execute the same scheduled operation, and you do not want to overload the server by starting them all simultaneously. The randomization function is not available when the schedule mode is server prompted.

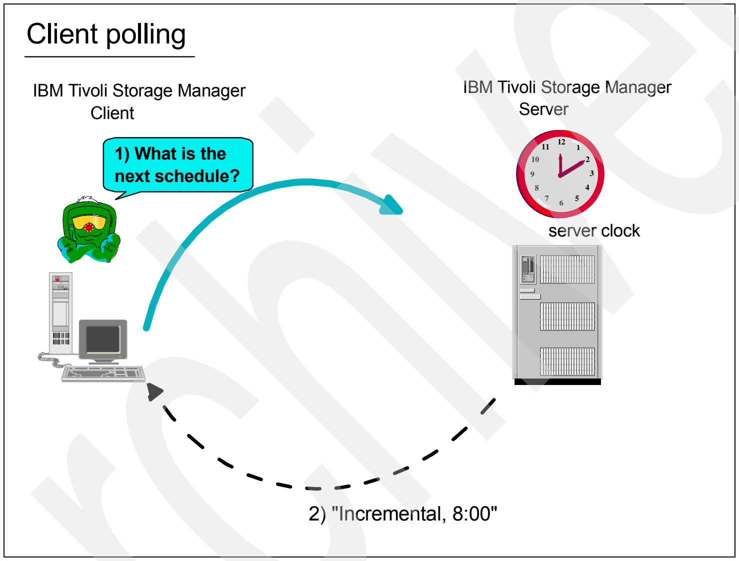

Figure 6-3 Client polling scheduling

Server prompted

If a client has selected the server-prompted mode, the client first contacts the server to notify the server that it is running, and the client is able to receive new schedule notifications. The server responds to let the client know whether or not it has schedules. The client then sleeps until the server contacts it. If the server detects that a new schedule is ready to run, it contacts the client and informs it that the client must start a new operation. The client receives the schedule definition and starts the operation.

Figure 6-4 shows an example of the communications between client and server when a scheduled operation must run on the client.

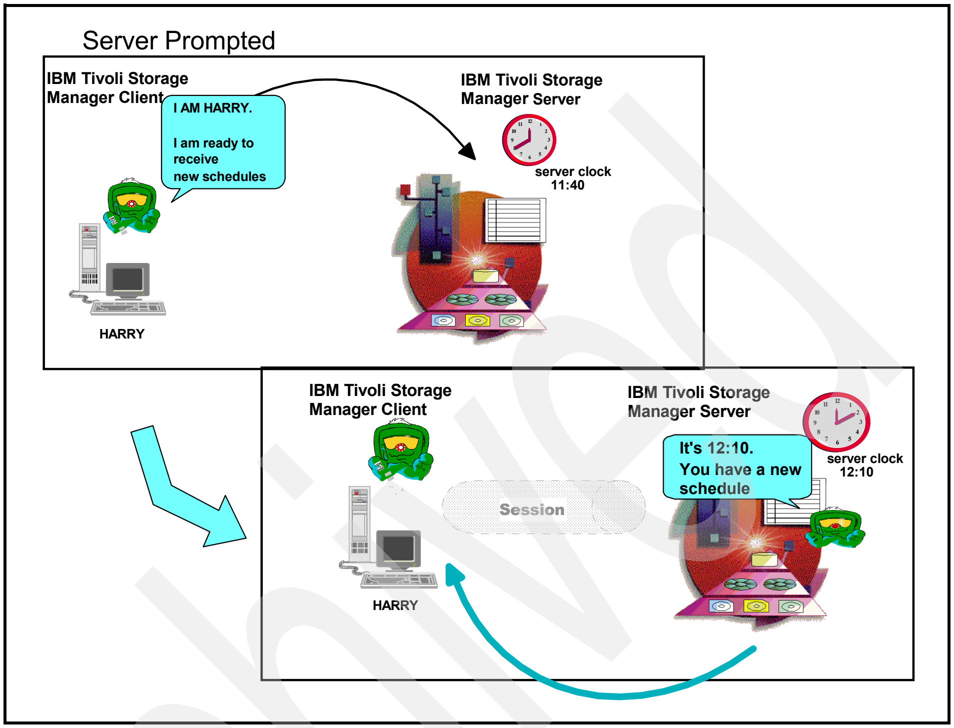

Figure 6-4 Server-prompted scheduling

If you use server-prompted scheduling, you can quickly and easily rerun a failed backup process by restarting it from the server. The other advantage of server-prompted scheduling is that you do not have the continued regular polling traffic from each client that is required by the client polling method (although this polling traffic is minimal).

Retry and randomization

A number of server parameters control the maximum number of concurrent client sessions allowed to connect to the Tivoli Storage Manager server, and some server parameters control the percentage of these sessions that are scheduled. If you restrict the number of scheduled sessions allowed on the server, you prevent a client from running a schedule when the maximum number of sessions has been reached. Through options that you can set globally at the server or individually for each client, the client can retry a certain number of times to run the schedule, with a specified time interval between retries.

As mentioned earlier, the server can randomize the start time of a client schedule within the configured startup window duration for clients in client-polling mode. The randomization parameter, controlled with SET RANDOMIZE, specifies the percentage of the startup window in which random start times are calculated. For example, if RANDOMIZE is set to 50% and a schedule has a startup window duration of two hours, a client in polling mode can be assigned a start time anywhere within the first hour (50% of two hours).

The retry and randomization options provide considerable flexibility in balancing the network load.

Consider a scenario where an administrator associates 100 workstations with a backup schedule that has a startup window between 2 a.m. and 6 a.m. every Friday. If the randomization option is set to 50%, Tivoli Storage Manager staggers the start times of the 100 backup sessions so that they start at different times between 2 a.m. and 4 a.m. The randomization prevents a large bottleneck from occurring if the 100 schedules all start at 2 a.m. Note that the more clients associated with a schedule, the larger the startup window must be in order for the randomization to be effective.

Logging scheduled events

Scheduled operations, also referred to as *events*, are stored in the Tivoli Storage Manager database. The results of scheduled events are stored in a log (which is really a database table). You can find other information regarding completed schedules in the server activity log.

You can view which schedules ran successfully, which were missed, and which are scheduled to run in the future. If you are interested in only those schedules that failed or were missed, you can view only those schedules out of the many that ran.

Detailed reports of the schedule are logged at the client level, with high-level data being sent to the server for logging (such as completion status, number of bytes sent, and so on). High-level results are sent to the event log so you can view whether or not the schedules ran successfully.

A number of server parameters configure how long event and schedule data is kept in the event and activity logs. The default for event data is 10 days, while the default for the activity log is one day. You also can manage the activity log by size rather than date if so desired.

6.3.2 Tivoli Storage Manager reporting

As with any application, a Tivoli Storage Manager environment includes a series of tasks that must be performed regularly. To perform these tasks in a timely and efficient way requires a set of resources, such as space in storage pools or in the Tivoli Storage Manager database, and tape drives and volumes. Scheduled operations must complete in a timely manner and without failures. In large environments, the number of operations can be high, and managing them effectively can be complex.

Reports are useful for verifying that the tasks set up for Tivoli Storage Manager to perform are carried out in a timely and efficient way.

Required reports

Reports and their content are very subjective; they depend on the installation's requirements and on each administrator's personal approach to Tivoli Storage Manager. One of the first things for an administrator to determine is the kind of operations Tivoli Storage Manager performs in each installation. For example, is Tivoli Storage Manager used only to back up files, or is Tivoli Storage Manager for Space Management also used? Is data being archived and are disaster protection procedures in place? The answers to these and other questions help build a set of reports that is appropriate to specific installation requirements.

Describing the exact commands required to generate these reports is beyond the scope of this book; however, the reports can all be generated using standard Tivoli Storage Manager administrative commands and SQL queries. Information about reporting queries is available in the companion book, *IBM Tivoli Storage Manager Implementation Guide*, SG24-5416.

Daily summary report

The Daily summary report or overview report is the most basic type of report. This report shows whether Tivoli Storage Manager is doing all that is expected and whether it summarizes all failed operations, such as missed backups or server errors. The Daily summary report must be as short as possible, to facilitate reading and understanding the information. It can include such information as the following list of all scheduled events that failed:

- Files that failed during backup
- Summary information about the amount of data transferred to the Tivoli Storage Manager server
- Space usage trends on the Tivoli Storage Manager server
- Removable media (tape) errors
- Server error messages

All information in the Daily summary report relates to the last 24-hour period. You can also generate a report once a week, in which case, of course, it must contain a summary of all activities and problems that occurred during the week.

Detail reports

Detail reports help explain a particular aspect of the Tivoli Storage Manager environment. The many possible Detail reports can be classified into the following categories:

- Client activity and traffic reports detail the activities performed on individual clients or groups of Tivoli Storage Manager clients. They can show the amount of data transferred to and from the server, categorized by activity type, such as backup, archive, or HSM. These reports are helpful in identifying delays in client activity and the reasons for the delays.
- Server background processes show the server activities that manage the data arriving from individual clients. Examples of these activities are migration, reclamation, and storage pool backup. It is useful to know when these activities run and their duration. You can use this information to estimate the impact of new workloads or server processes on current server activities.
- Server storage pool space utilization reports show server storage pool space utilization by the total storage pool space, client node, client file space, or storage pool. These reports can be point-in-time or trend reports. They are useful in determining the amount of storage media required to accommodate current or new workloads.
- Server database utilization reports monitor the Tivoli Storage Manager server database and log utilization to ensure that adequate space is always available. Tivoli Storage Manager does not work properly if the database or log fills up, so this situation must be avoided.
- Backup/archive schedules show scheduled client node operations that use the Tivoli Storage Manager scheduler. The reports show scheduled events that failed, which is the basic exception report you monitor, and all scheduled events that occurred. Events scheduled using utilities, such as a UNIX cron job or a third-party scheduling package, are not logged on the Tivoli Storage Manager server, so another reporting mechanism must be used.
- Administrative schedules can be used to automate administrative operations. Tivoli Storage Manager reports an administrative event as completed when the scheduled operation is started and not when it has been completed. Therefore, you want to report on the completion status of administrative scheduled events.

- Server configuration reports show Tivoli Storage Manager server configuration parameters such as management. These reports can be used as references when analyzing other reports.

Server information storage

Most of the information used to create reports is stored on the Tivoli Storage Manager server, but some of it is stored only on client nodes.

Information about the server

The Tivoli Storage Manager server uses three main sources of information for creating reports:

- Database

  This source is the most important. The Tivoli Storage Manager database contains all server and most client definitions. It is the prime source of static information. When database information is requested and given, it is in the form of snapshot or point-in-time data; trends cannot be seen.

- Activity log

  This report contains all server messages for the past several days. Entries are kept for a configurable number of days. Messages are logged for client sessions, scheduled operations, automated server processes, and any errors that might occur. The activity log displays historical information: You can review the progress of operations over time.

  Client messages can be logged as events in the activity log. Commands are available to select all, none, or a subset of events to be logged.

- Accounting log

  Tivoli Storage Manager accounting can be optionally activated. Once started, records are automatically collected, and the log is written to the dsmaccnt.log file in the server installation directory. The log file contains a record for each client session that terminates. The record consists of comma-separated text and numeric values, so the log can be conveniently loaded into a spreadsheet or another data analysis package. It offers summary information for the activities performed during a session, such as files and bytes backed up, archived, or migrated, session wait times, and total data transferred.

Information about the client node

The Tivoli Storage Manager client has the following two main sources of information:

- Client error log

  Tivoli Storage Manager writes error information to the dsmerror.log file, usually for situations where the client did not succeed in contacting the server. By default, this log is written in the Tivoli Storage Manager client installation directory; however, you can change the location by setting a client option.

- Scheduler log

  The dsmsched.log file contains information for all scheduled operations, information such as the name of files that are backed up or archived, failures and errors, and backup summary statistics. By default the scheduler log file is written to the current directory where the client scheduler is started; however, you can change the location by setting a client option.

Tivoli Storage Manager client sessions that are not started with the scheduler (for example, GUI or command-line sessions) do not write output to a file by default. The output is lost unless the messages are saved in some way, such as by redirecting standard output and errors to a file.

Error and summary information is propagated from the client to the server activity log. Client error and information messages are identified in the server activity log by the prefix characters ANE. This information is stored for both scheduled and unscheduled Tivoli Storage Manager client sessions.

Central error logging

Tivoli Storage Manager events can be centrally logged, monitored, and reported using industry-standard interfaces. Central error logging means the implementation can be integrated with system management applications, providing centralized control.

Certain Tivoli Storage Manager client messages can be logged as events on the server. Client messages can be collected in one central point. The intent of client message logging is to log problems encountered during a Tivoli Storage Manager client operation; therefore, only messages indicating an error condition are logged as events. The only exception is client backup statistics, which also can be centrally logged.

All events that are to be logged must be enabled by either message number or severity. Enabled client events that are logged to the Tivoli Storage Manager server are, by default, stored in the activity log and displayed on the server console.

Client and server event reporting

IBM Tivoli Storage Manager can send client and server events to external interfaces, allowing Tivoli Storage Manager to integrate with other system management packages. Supported interfaces include SNMP managers such as NetView® for AIX, CA Unicenter, and HP OpenView; other available interfaces are to Tivoli Enterprise Console, NetView for MVS, the Microsoft Windows event log, a user-written exit, and writing direct to a file. Interfaces that receive event data are called *event receivers*. Each event message, whether client or server, can be enabled for any of the supported receivers. It is possible to enable one message or a group of messages for more than one receiver. As with client event logging, events are enabled for receivers by message number or severity.

SNMP server heartbeat monitoring

SNMP is also used to monitor network elements from a central point. It enables the monitored systems to send traps notifying the SNMP manager about events taking place on the local system. As well as sending traps, a heartbeat monitor can monitor whether managed Tivoli Storage Manager servers are still alive.

To take advantage of SNMP monitoring, Tivoli Storage Manager includes an interface for SNMP, which is distributed with the server as an SNMP subagent. It is supported for servers running on AIX, HP-UX, Solaris, Linux, and Microsoft Windows. Communication between the server and the SNMP manager is established through one of the following connection channels:

- Tivoli Storage Manager server <—> SNMP subagent <—> SNMP agent <—> SNMP manager
- Tivoli Storage Manager server <—> SNMP agent <—> SNMP manager

To enable communication between SNMP subagent and SNMP agent, the SNMP agent must support the Distributed Protocol Interface (DPI).

SQL queries and ODBC interface

Tivoli Storage Manager provides an SQL interface that supports queries to its internal database. The interface is read-only and includes a SELECT command, which can be used on any server platform, and an Open Database Connectivity (ODBC) driver for the Windows backup/archive client.

SELECT command

The SQL interface represents Tivoli Storage Manager information in the form of relational tables containing rows and columns that can be accessed by the SELECT command. The SELECT command uses standard SQL syntax compliant with the SQL92/93 standard and can be used only on the administrative command-line client.

Because SQL processing uses database resources, long-running or very complicated SELECT statements can slow server performance significantly. Therefore, resource-intensive queries display a confirmation message, offering the possibility of aborting the query before executing it.

ODBC driver

ODBC is a standard interface between SQL database engines and front-end applications. You can use products such as Microsoft Access to graphically construct SQL queries, which are then dispatched to the database (in this case, the Tivoli Storage Manager database). The SELECT statement results are returned in tabular form and can be processed to be displayed as charts or tables. The Tivoli Storage Manager ODBC driver is shipped only with the Windows backup/archive client package.

Operational reporting

This section discusses operational reporting, a feature made available with Tivoli Storage Manager V5.2.2. Operational reporting is available with Tivoli Storage Manager for Windows. If you do not run a Tivoli Storage Manager Windows server, operational reporting also is available as a separately installable package for Windows. Operational reporting can report on Tivoli Storage Manager servers on any supported platform.

Current complementary products include several items for generating reports on Tivoli Storage Manager. These reports, created for long-term analysis, provide a mechanism for making Tivoli Storage Management events, performance, and fulfilment of business requirements readily visible using a variety of formats and reporting levels. The reports are appropriate and informative for information systems technicians, Tivoli Storage Management administrators, and company executives. The reports usually include different views and diagrams to show how the Tivoli Storage Manager environment has developed during the analysis time frame.

Establishing such complex reports requires a lot of data, which comes from the Tivoli Storage Manager database. The database itself was not specifically developed for data warehousing and online analytical processing, so the required records are first transferred to an external relational database management system (RDBMS), which generates the desired reports. This process can be very time consuming, so it usually is executed only once a day. Because of their emphasis in the long term, the reports do not include real-time information about the current state of a Tivoli Storage Manager server.

Operational reporting, alternatively, is made to support Tivoli Storage Manager administrators in their daily work. You usually execute a `query` or SELECT command to discover current issues on a Tivoli Storage Manager server. These tasks can be repeated daily or hourly, displaying the most current data to keep Tivoli Storage Manager running smoothly and resolve any issues as quickly and easily as possible.

Operational reporting views a Tivoli Storage Manager server as being in one of two states, Running Smoothly or Needs Attention, determined by customizable rules. This information is automatically determined and sent in the subject line of an e-mail. The e-mail provides access to a customizable status report. If a server needs attention, the e-mail also describes the issues and provides recommendations on how to get Tivoli Storage Manager running smoothly again.

For example, you can configure Tivoli Storage Manager operational reporting to query Tivoli Storage Manager servers every day at 5:00 a.m., generate a report with the designated information, and send the report as an e-mail with a subject line indicating whether the server is running smoothly or needs attention. A server requires attention if any of the outstanding issues are specified in the report, issues such as the failure of more than one client schedule or the detection of fewer than 10 scratch volumes. You can also provide recommendations on what to do when issues arise.

If the report indicates that a server needs attention, it highlights the issues and provides recommendations for resolving them. Operational reporting also includes a special kind of report, an operational monitor, which you can configure to run hourly to check for issues. Whereas a report is sent regardless of any issues, a monitor notifies you only if any issues arise. You are notified by e-mail or a message is sent to your Microsoft Windows desktop. You can customize and share XML-based operational report and monitor templates with others.

Tivoli Storage Manager operational reporting enables the following benefits:

- Reduces the amount of time required to administer Tivoli Storage Manager
- Provides customizable daily operational summary reports
- Provides customizable hourly monitors

- Supports multiple Tivoli Storage Manager servers (any version, any platform)
- Supports multiple reports and monitors per server
- Produces reports that can be viewed interactively or from a Web site
- Provides quick status identification using color highlighting
- Allows sharing of custom report and monitor templates
- Identifies issues and provides recommendations
- Notifies administrators:
  - Automatically by e-mail with the system status in the subject line
  - By messages sent to administrator's desktop
- Notifies clients of failed or missed schedules
- Is easy to use because:
  - It runs as a Microsoft Windows service.
  - It is integrated into the Tivoli Storage Manager management console by a snap-in.
  - Defaults are provided out of the box.
  - A small number of settings must be configured with a simple user interface.

Example report output

This section describes sample outputs from Tivoli Storage Manager operational reporting. These examples demonstrate the previously described features and how they keep Tivoli Storage Manager servers running smoothly to reduce administration time and effort.

Web summary page

You can easily configure operational reporting to generate reports accessible through a Web browser using a local Web server.

1. From your management console, right-click **Tivoli Storage Manager** and select **TSM Operational Reporting** as shown in Figure 6-5 on page 314.

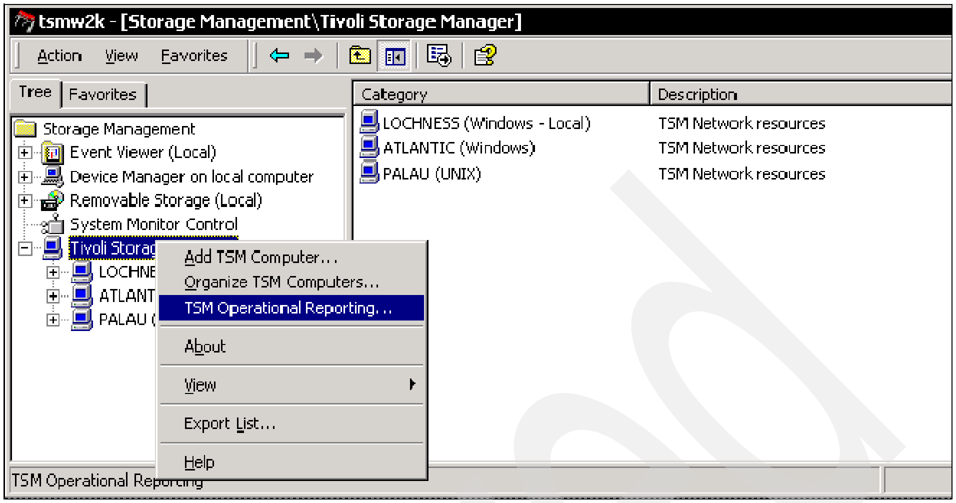

Figure 6-5 Configure Web summary reports

2. Select the **Summary Information** tab on the next panel. You must configure a URL to access the reports, provide the name of the directory from which the reports can be accessed through the local Web server, and provide the default file name for the index or default page, as shown in Figure 6-6.

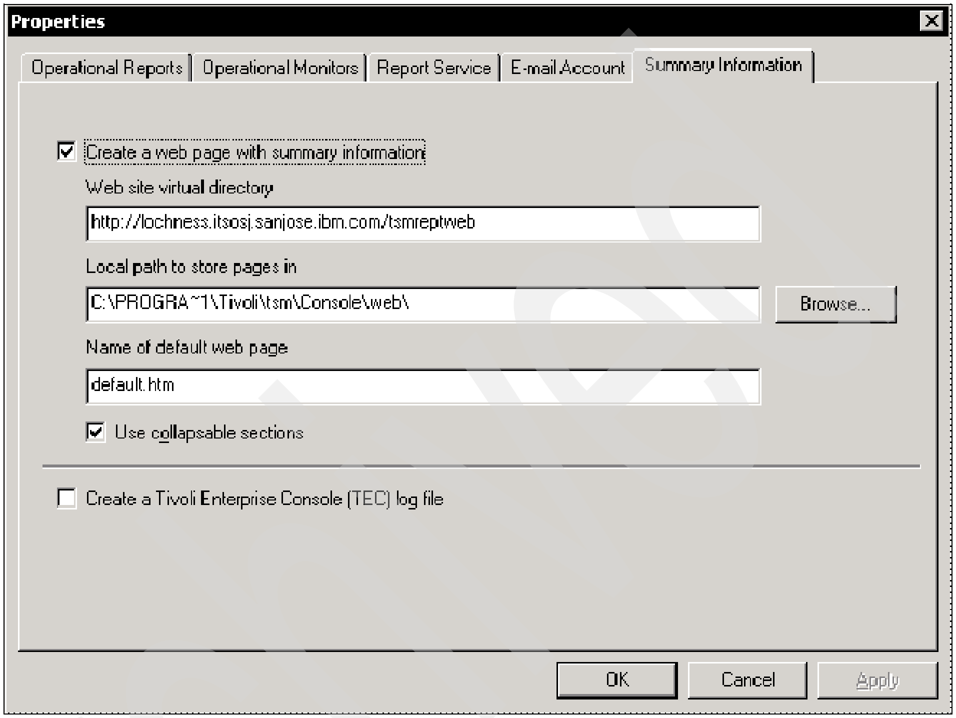

Figure 6-6 Web summary reports configuration

If you enable the Web summary feature, operational reporting generates a Web page that combines references and links to all of the latest generated monitors and reports, as shown in Figure 6-7. By clicking the links you can reach specific reports or monitors.

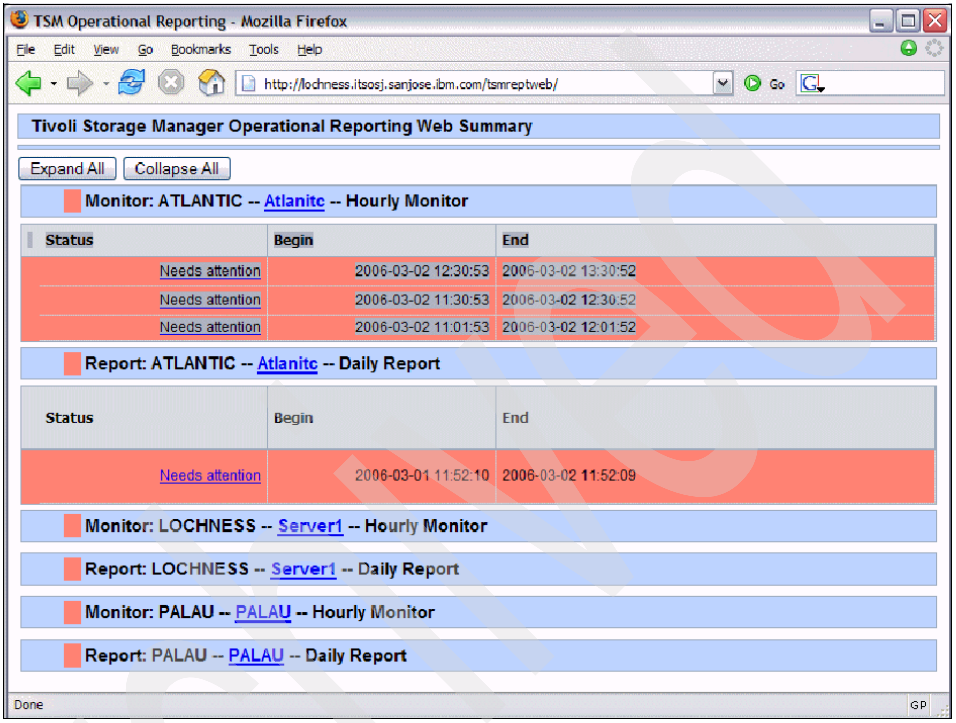

Figure 6-7 Operational report Web summary page

Hourly monitor

You can create monitors to observe only the main vital parameters of a Tivoli Storage Manager server. This monitor is usually refreshed every hour. If an issue occurs, the monitor results are forwarded to a defined e-mail account or to the Web summary page, as shown in Figure 6-8.

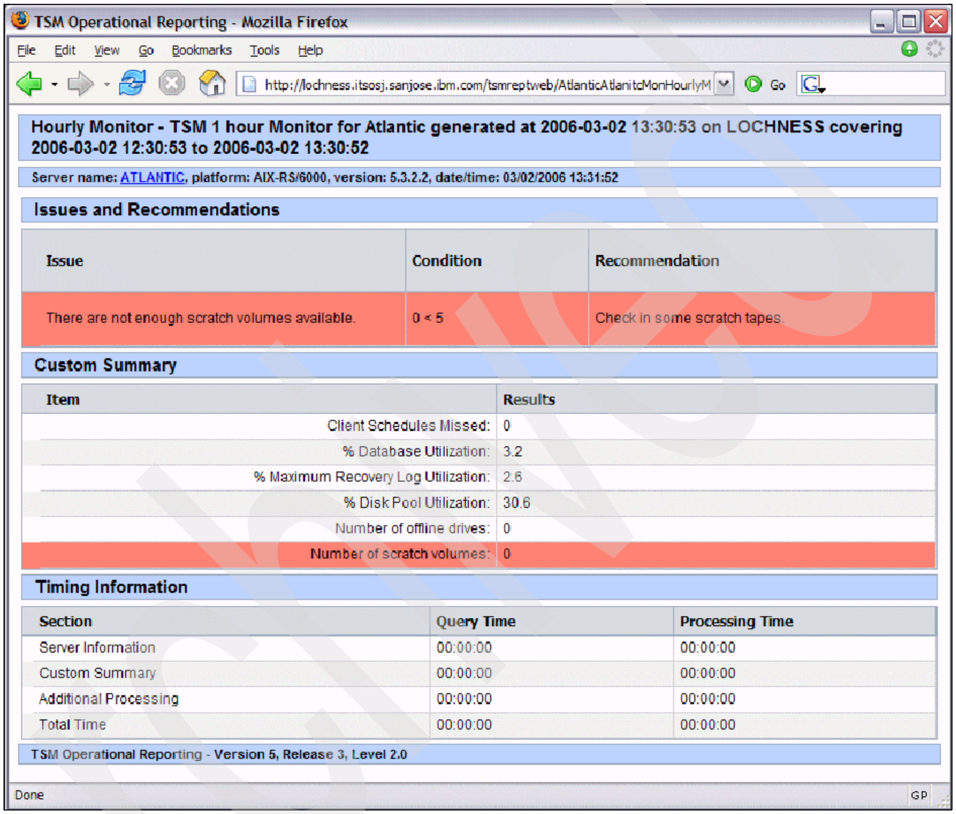

Figure 6-8 Operational reporting hourly monitor

Daily report

The daily report is a summary of all activities that took place during the previous 24 hours. It includes useful information about client and server activities, such as execution results of schedules, client processes, missed files, throughput, bytes transferred, and server status. Furthermore, it includes summaries for specific server processes such as migration, reclamation, and tape mounts across the time line. The bar charts shown in Figure 6-9 and Figure 6-10 on page 319 are examples of such summaries.

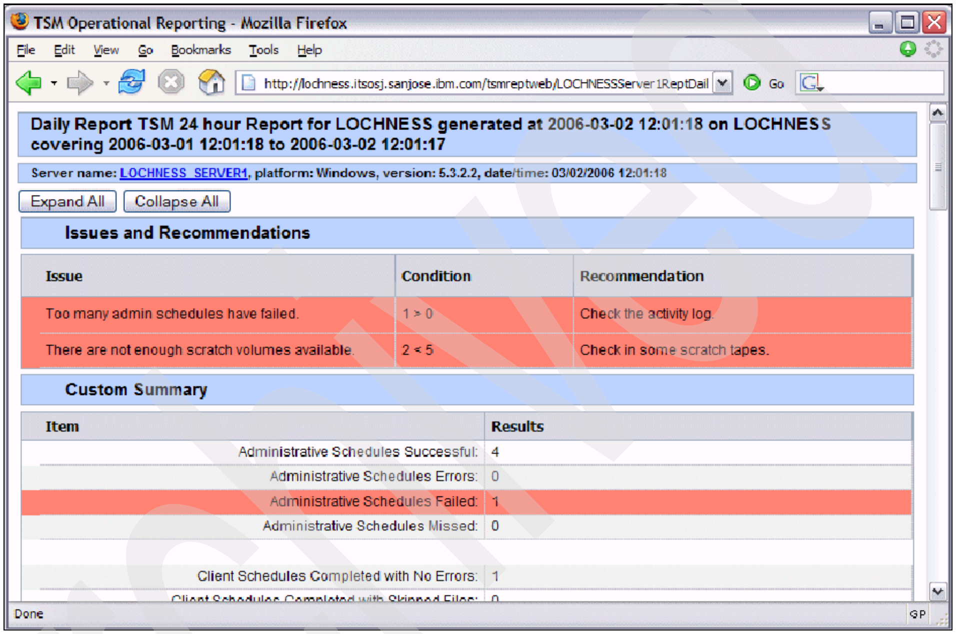

Figure 6-9 Operational reporting daily report summary

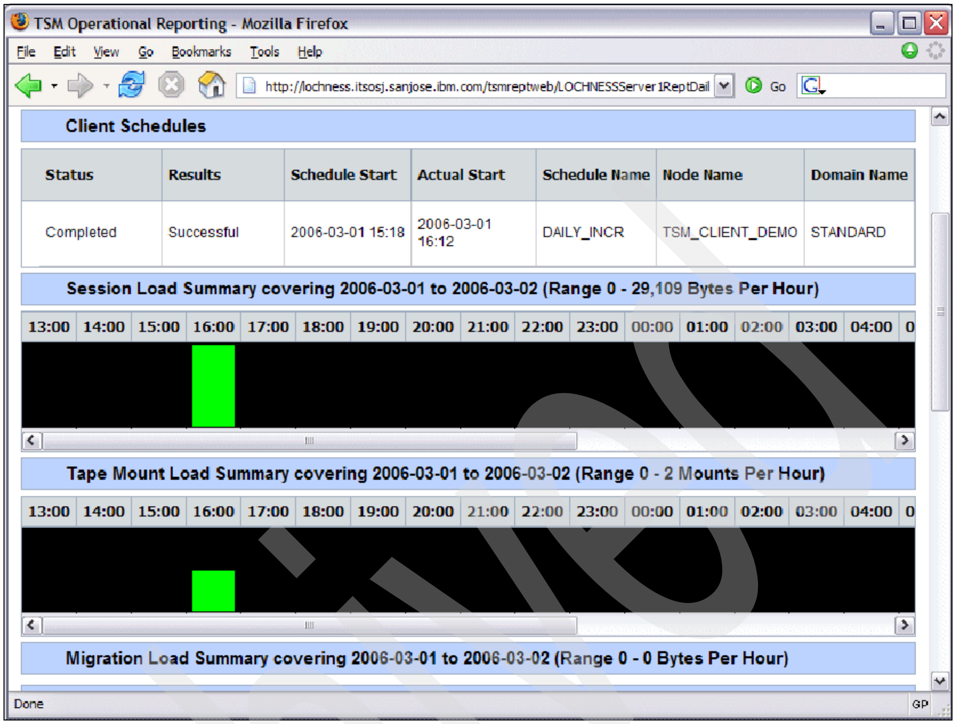

Figure 6-10 Operational reporting daily report load summary

E-mail notifications

Tivoli Storage Manager operational reporting can notify you by e-mail as soon as issues are detected. You can forward the defined reports and monitors using this feature. These e-mails can include the issue list in plain text or HTML format, or they can forward the address of the Web summary page, which you can use to access the latest monitor or report.

Tivoli Storage Manager operational reporting can notify the dedicated contact persons of defined client nodes by e-mail when an associated schedule has been missed or has failed (see Example 6-1). This feature simplifies the process of informing you of possibly misconfigured Tivoli Storage Manager clients.

Example 6-1 E-mail notification for missed schedule

```
To: catfish@aquarium.com
cc:
Subject: Missed TSM Schedule for node JAMAICA.

Hello Mr. Catfish,
```

```
You are receiving this automatic notification message because you are listed as
the contact for node JAMAICA on TSM server CLYDE.

The node has missed its scheduled backup.

Typical reasons for missed schedules are:

  The computer is not on the network.
  The scheduler is not installed.
  The scheduler is not running because it is not set to start at boot time.
  The scheduler is not running because it had an error.
  The scheduler's option file is not pointing to the right server.

It may help to review the *.log files in the baclient directory to identify the
problem.

Please reply if you:

  [] No longer want to be automatically notified of missed schedules.
  [] No longer want to be associated with this schedule.
  [] Need a different schedule -- specify your preferred date and times.
  [] Need help fixing the problem -- attach the client schedule and error logs
if possible.
  [] Other -- provide details.

Thank you.

This automatic message was sent from machine CLYDE.
```

This example describes the results Tivoli Storage Manager operational reporting can provide. This useful triggering and conditioning of Tivoli Storage Manager-related information helps reduce the daily amount of time and effort required to administer a Tivoli Storage Manager environment.

6.4 Administration Center monitoring and reporting

The Tivoli Storage Manager Administration Center provides monitoring and reporting capabilities. Our discussion in this section includes the following topics:

- 6.4.1, “Health monitor” on page 321
- 6.4.2, “Administration Center reporting” on page 323

6.4.1 Health monitor

Use the health monitor to determine the overall status of server operations and to obtain detailed information about:

- Scheduled events
- Database and recovery logs
- Activity log
- Storage device status

The health monitor must be configured only once, regardless of how many administrators and servers you define to the Administration Center. The configuration applies to all servers. Figure 6-11 shows the health monitor start page, where you can select which server you want to monitor. Note the message at the bottom of the panel: The health monitor is not an ad hoc picture but is updated frequently, based on the refresh interval.

> **Note:** The refresh interval countdown displayed in the health monitor Web page is not intended to match the actual cycle the health monitor is using to contact servers. Instead, the refresh interval countdown is when the Web page itself is scheduled to refresh.

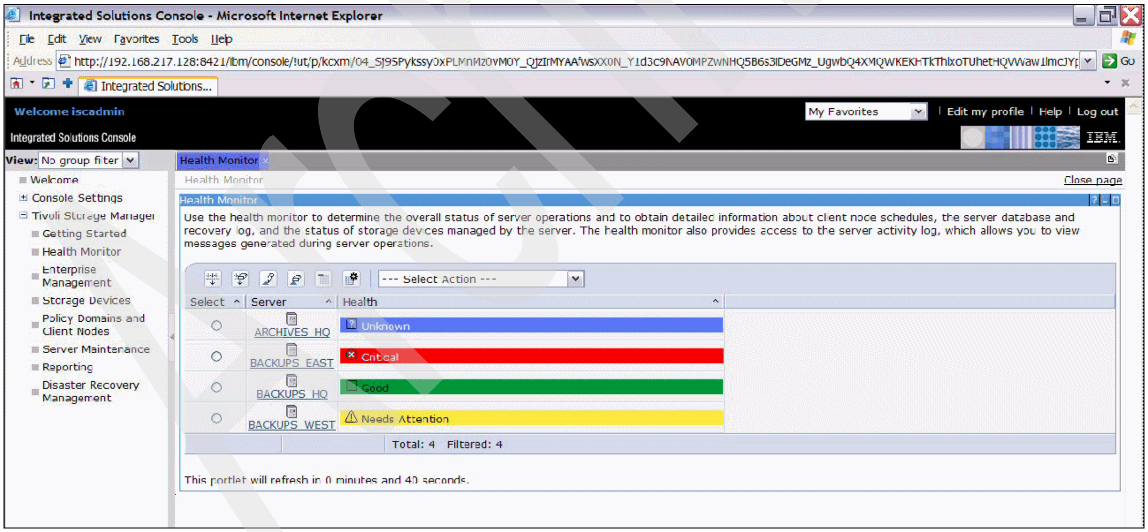

Figure 6-11 Health monitor

Figure 6-12 shows the health monitor report for a specific server. You can request the detailed database information being displayed. The top of the work area shows a summary that indicates areas that might require your attention. You can expand each of these areas for more details in the panels below.

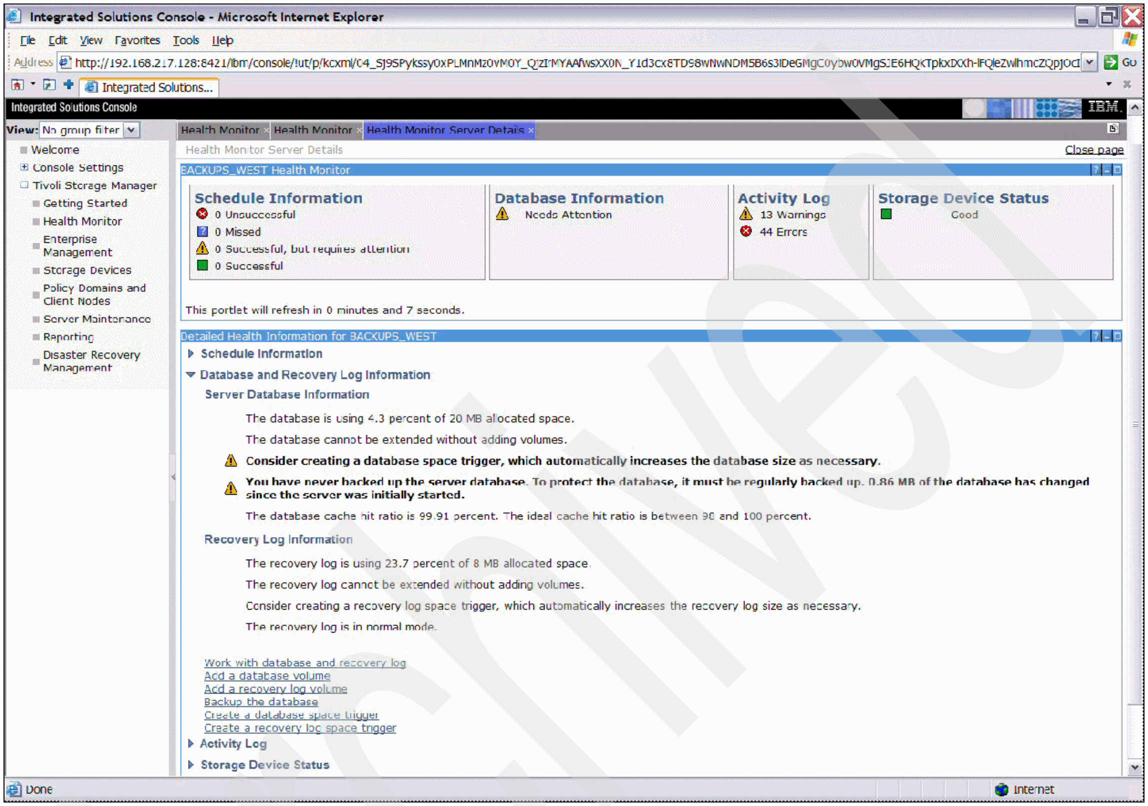

Figure 6-12 Health monitor, database information details

Use Figure 6-13 on page 323 as an example of how the different health states are displayed. The health monitor determines the health status (shown in Figure 6-11 on page 321, for example, Good or Needs Attention) by evaluating a set of predefined fixed rules. No single condition causes the change in status, but rather the health monitor looks at a variety of items to determine the overall status.

For detailed information about the health monitor and how it works, see *IBM Tivoli Storage Manager Problem Determination Guide Version 5.3.2*, SC32-9103.

6.4.2 Administration Center reporting

The reporting task in the Administration Center provides access to a set of predefined reports. The following reports are currently available:

- Usage reports
- Security reports

From the reporting window, shown in Figure 6-13, you can choose which servers to report on.

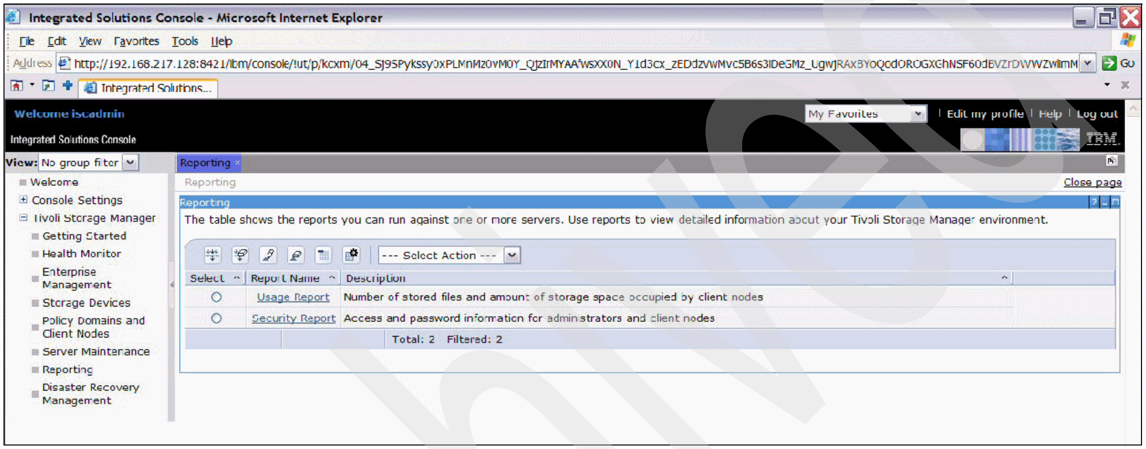

Figure 6-13 ISC reporting

Figure 6-14 shows a sample usage report. You can collapse sections of the report, and you can select rows by clicking the small arrow icon on the column or using the sort button.

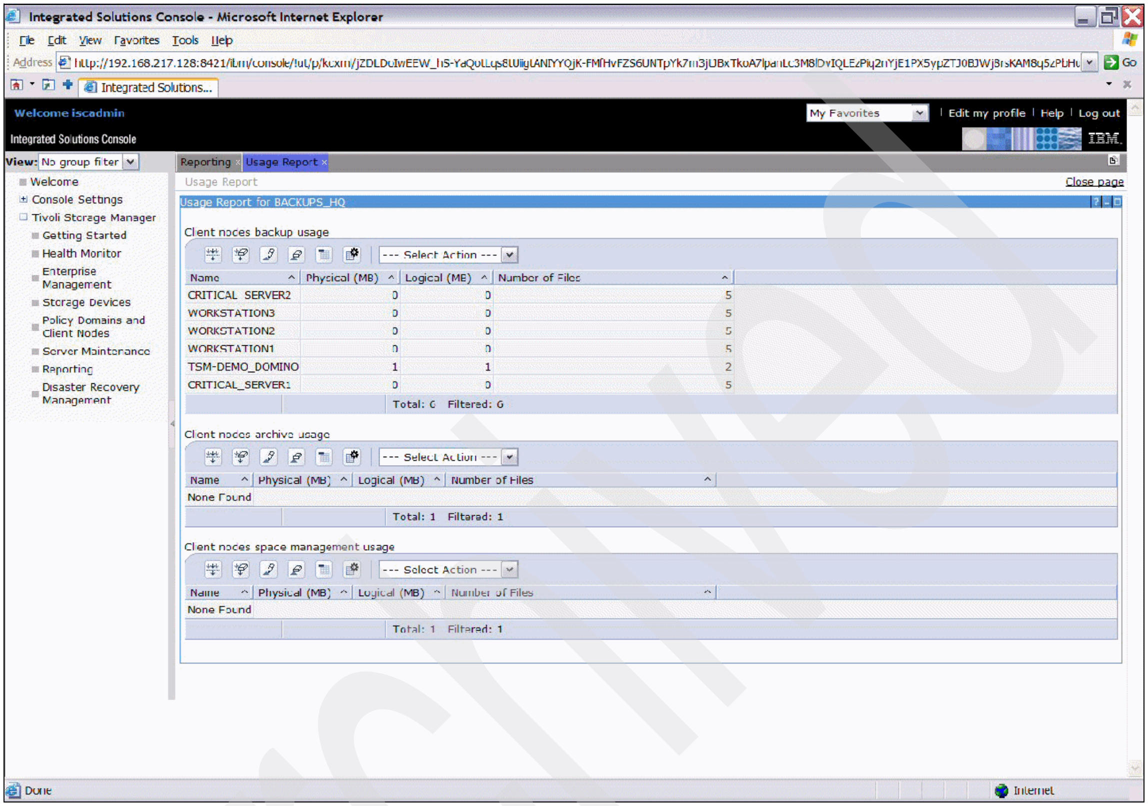

Figure 6-14 ISC reporting, usage report

The rich filter and sort functions of ISC reporting makes it easy for you to access the information you are looking for.

6.5 Summary

This chapter discussed miscellaneous tips and tricks for managing your Tivoli Storage Manager environment.

Part 3

Appendix

The appendix for this deployment guide contains planning tables. See Appendix A, "Worksheets for planning and sizing" on page 327.

A

Worksheets for planning and sizing

This appendix includes the following worksheets you can use in planning of an IBM Tivoli Storage Manager implementation:

Client requirements

Use Table A-1 as a checklist of client requirements.

Table A-1 Client requirements worksheet

Category	Client 1	Client 2	Client 3	Client 4
Client name				
Contact information				
Operating system				
Total storage available (GB)				
Total storage used (GB)				
GB changed per backup				
Number of files backed up				
Data compression				
Backup window times				
Backup number of hours				
Required recovery time				
IBM Tivoli Storage Manager recovery time				
GB copied per archive				
Number of files archived				
Number of archives kept				
Archive frequency				
Archive window times				
Archive number of hours				
Number of image backups				
Image backup frequency				
Number of backup sets				
Backup set frequency				

Category	Client 1	Client 2	Client 3	Client 4
Policy domain				
Client option set				

Storage policy requirements

You can use Table A-2 to list storage policy requirements.

Table A-2 Storage policy requirements worksheet

Category	Example 1	Example 2	Example 3
Group name			
Number of backup versions			
Backup file retention period			
Number of deleted versions			
Last deleted file version retention period			
Archive retention period			
Off-site copies			
On-site collocation			
Off-site collocation			
Image backup retention			
Backupset retention			

Database

Table A-3 is a database worksheet.

Table A-3 Database worksheet

Database volume	Filename (primary)	Size (MB)	Filename (copy)	Size (MB)

Recovery log

Keep track of logs with Table A-4.

Table A-4 Recovery log worksheet

Log Volume	Filename (primary)	Size (MB)	Filename (copy)	Size (MB)
	Total		Total	

Device configuration and volume history

Table A-5 is a worksheet for listing device configuration and volume history.

Table A-5 Device configuration and volume history worksheet

Category	Size (MB)

Category	Size (MB)
Total	

IBM Tivoli Storage Manager disk requirements

Use Table A-6 to list Tivoli Storage Manager disk requirements.

Table A-6 Total IBM Tivoli Storage Manager disk requirements worksheet

Category	Size (MB)
IBM Tivoli Storage Manager software (dependent on platform)	
IBM Tivoli Storage Manager database	
IBM Tivoli Storage Manager recovery log	
IBM Tivoli Storage Manager primary storage pools	
Device configuration table and volume history table	
Other (RAID, operating system)	
Total	

Tape drive configuration

Table A-7 is a worksheet for noting tape drive configuration information.

Table A-7 Tape drive configuration worksheet

Category	Option
Library model	
Number of drives	
Drive model	
Number of on-site tape volumes	
Number of off-site tape volumes	
Number of database volumes	
Number of scratch tapes	
Number of backup set tape volumes	
Total tape volumes required	

Administrator ID

You can use Table A-8 on page 333 to list administrator IDs.

Table 6-1 Administrator IDs worksheet

Functions	IBM Tivoli Storage Manager ID	Authority

Functions	IBM Tivoli Storage Manager ID	Authority

License requirements

Table A-8 is a worksheet for listing license requirements.

Table A-8 License requirements worksheet

Category	Required
Server type	
Client connections	
Network connections	
Open systems environment clients	
Space management	

Abbreviations and acronyms

AAS	Advanced Administrative System
ACL	access control list
ACSLS	Automated Cartridge System Library Software
ADSM	AdStar Storage Manager
AES	Advanced Encryption Standard
AIX	Advanced Interactive eXecutive
AME	application-managed encryption
API	application programming interface
BCP	business continuity planning
BRP	business recovery plan
CAD	client acceptor daemon
CBMR	Cristie Bare Metal Restore
CDP	Continuous Data Protection for Files
CFI	Complete File Index
CIFS	Common Internet File System
CIM	Common Information Model
CLI	command-line interface
COOP	continuity of operations plan
DB2	Database 2™
DES	Data Encryption Standard
DLT	Digital Linear Tape
DMA	data management application
DMZ	demilitarized zone
DNS	Domain Name Service
DPI	Distributed Protocol Interface
DRM	Disaster Recovery Manager
DRP	disaster recovery plan
DVD	Digital Video Disk
DWDM	Dense Wavelength Division Multiplexing
EFS	Encrypted File System
EKM	Encryption Key Manager or Enterprise Key Manager
ERP	enterprise resource planning
ESS	Enterprise Storage Server
FAT	file allocation tables
FTP	File Transfer Protocol
GPFS	General Parallel File System
HACMP	High-Availability Cluster Multi-Processing
HBA	host bus adapter
HSM	hierarchical storage manager
HTTP	Hypertext Transfer Protocol
HTTPS	HTTP Secure
HVD	Holographic Versatile Disc
IBM	International Business Machines Corp.
IPV6	Internet Protocol Version 6
IRP	incident response plan
ISC	Integrated Solutions Console
ISO	International Standards Organization
ITIL	Information Technology Infrastructure Library
ITSO	International Technical Support Organization
JFS	Journaled File System
JRE	Java Runtime Environment
LAN	local area network
LME	Library Managed Encryption
LTO	Linear Tape Open

LVM	Logical Volume Manager
LVSA	Logical Volume Snapshot Agent
MIB	management information base
MMC	Microsoft Management Console
MOM	Microsoft Operations Manager
MOSS	Microsoft Office SharePoint Server
MSCS	Microsoft Cluster Server
MTS	Microsoft Transaction Server
MTU	Maximum Transmission Unit
MVS	Multiple Virtual Storage
NAS	network attached storage
NCS	Novell Cluster Services
NDMP	Network Data Management Protocol
NFS	network file system
NTFS	NT File System
ODBC	Open Database Connectivity
PCI	Peripheral Component Interconnect
PDF	Portable Document Format
PIT	point in time
POSIX	Portable Operating System Interface
RAID	Redundant Array of Independent Disks
RAM	random access memory
RDBMS	Relational Database Management System
RHEL	Red Hat Enterprise Linux
SAN	storage area network
SATA	Serial Advanced Technology Attachment
SCSI	Small Computer System Interface
SDK	Software Developer's Kit
SLA	service-level agreement
SLES	SUSE Linux Enterprise Server
SME	system-managed encryption
SMIT	System Management Interface Tool
SNMP	Simple Network Management Protocol
SQL	Structured Query Language
SSL	Secure Sockets Layer
SVC	SAN Volume Controller
TCO	Cost of Ownership
TCP/IP	Transmission Control Protocol/Internet Protocol
TPC	TotalStorage Productivity Manager
UDB	Universal Database
URL	Universal Resource Locator
VCB	VMware Consolidated Backup
VCS	Veritas Cluster Server
VOLSER	Volume Serial
VSS	Volume Shadow Copy Service
VTL	Virtual Tape Libraries
WAN	wide area network
WPAR	workload partition

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

For information about ordering these publications, see “How to get Redbooks” on page 341. Note that some of the documents referenced here may be available in softcopy only.

- *IBM Tivoli Storage Manager Versions 5.4 and 5.5 Technical Guide*, SG24-7447
- *Backing up Microsoft SQL Server with IBM Tivoli Storage Manager*, SG24-6148
- *Deployment Guide Series: IBM Tivoli Storage Manager Express*, SG24-7033
- *IBM System Storage N series*, SG24-7129
- *IBM Tivoli Storage Manager: Building a Secure Environment*, SG24-7505
- *IBM Tivoli Storage Manager for Advanced Copy Services*, SG24-7474
- *Using the IBM System Storage N Series with IBM Tivoli Storage Manager*, SG24-7243
- *IBM Tivoli Storage Manager in a Clustered Environment*, SG24-6679
- *IBM System Storage Business Continuity Solutions Overview*, SG24-6684
- *IBM System Storage Business Continuity: Part 1 Planning Guide*, SG24-6547
- *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548
- *Disaster Recovery Strategies with Tivoli Storage Management*, SG24-6844
- *IBM Tivoli Storage Management Concepts*, SG24-4877
- *IBM Tivoli Storage Manager Version 5.3 Technical Guide*, SG24-6638
- *IBM Tivoli Storage Manager: A Technical Introduction*, REDP-0044

Other publications

These publications are also relevant as further information sources:

- Tivoli Storage Manager server publications
  - *IBM Tivoli Storage Manager Messages*, SC32-0140
  - *IBM Tivoli Storage Manager Performance and Tuning Guide*, SC32-0141
  - *IBM Tivoli Storage Manager Problem Determination Guide Version 5.5*, SC32-0142
  - *IBM Tivoli Storage Manager for AIX Installation Guide Version 5.5*, GC23-5969
  - *IBM Tivoli Storage Manager for AIX Administrator's Guide Version 5.5*, SC32-0117
  - *IBM Tivoli Storage Manager for AIX Administrator's Reference Version 5.5*, SC32-0123
  - *IBM Tivoli Storage Manager for HP-UX Installation Guide Version 5.5*, GC23-5970
  - *IBM Tivoli Storage Manager for HP-UX Administrator's Guide Version 5.5*, SC32-0118
  - *IBM Tivoli Storage Manager for HP-UX Administrator's Reference Version 5.5*, SC32-0124
  - *IBM Tivoli Storage Manager for Linux Installation Guide Version 5.5*, GC23-5971
  - *IBM Tivoli Storage Manager for Linux Administrator's Guide Version 5.5*, SC32-0119
  - *IBM Tivoli Storage Manager for Linux Administrator's Reference Version 5.5*, SC32-0125
  - *IBM Tivoli Storage Manager for Sun Solaris Installation Guide Version 5.5*, GC23-5972
  - *IBM Tivoli Storage Manager for Sun Solaris Administrator's Guide Version 5.5*, SC32-0120
  - *IBM Tivoli Storage Manager for Sun Solaris Administrator's Reference Version 5.5*, SC32-0126
  - *IBM Tivoli Storage Manager for Windows Installation Guide Version 5.5*, GC23-5973
  - *IBM Tivoli Storage Manager for Windows Administrator's Guide Version 5.5*, SC32-0121

– *IBM Tivoli Storage Manager for Windows Administrator's Reference Version 5.5*, SC32-0127
– *IBM Tivoli Storage Manager for System Backup and Recovery 6.1 Installation and User's Guide*, SC23-6543

- Tivoli Storage Manager storage agent publications
  – *IBM Tivoli Storage Manager for AIX Storage Agent User's Guide Version 5.4*, SC32-0129
  – *IBM Tivoli Storage Manager for HP-UX Storage Agent User's Guide Version 5.4*, SC32-0130
  – *IBM Tivoli Storage Manager for Linux Storage Agent User's Guide Version 5.4*, SC32-0131
  – *IBM Tivoli Storage Manager for Sun Solaris Storage Agent User's Guide Version 5.4*, SC32-0132
  – *IBM Tivoli Storage Manager for Windows Storage Agent User's Guide Version 5.4*, SC32-0133
- Tivoli Storage Manager client publications
  – *IBM Tivoli Storage Manager for Macintosh: Backup-Archive Clients Installation and User Guide Version 5.5*, SC32-0143
  – *IBM Tivoli Storage Manager for NetWare: Backup-Archive Clients Installation and User Guide Version 5.5*, SC32-0144
  – *IBM Tivoli Storage Manager for UNIX and Linux: Backup-Archive Clients Installation and User Guide Version 5.5*, SC32-0145
  – *IBM Tivoli Storage Manager for Windows: Backup-Archive Clients Installation and User Guide Version 5.5*, SC32-0146
  – *IBM Tivoli Storage Manager for Space Management for UNIX: User Guide Version 5.5*, SC32-0148
  – *IBM Tivoli Storage Manager for HSM for Windows Administrator's Guide Version 5.5*, SC32-1773
  – *IBM Tivoli Storage Manager Using the Application Program Interface Version 5.5*, SC32-0147
- Tivoli Storage Manager Data Protection publications
  – *IBM Tivoli Storage Manager for Advanced Copy Services: Data Protection for Snapshot Devices for SAP Installation and User's Guide for Oracle Version 5.5*, SC33-8207
  – *IBM Tivoli Storage Manager for Advanced Copy Services: Data Protection for Snapshot Devices Installation and User's Guide Version for DB2 5.5.0*, SC33-8330

  - *IBM Tivoli Storage Manager for Application Servers 5.2: Data Protection for Application Server 5.0 Installation and User’s Guide*, SC32-9075
  - *IBM Tivoli Storage Manager for Informix Installation and User’s Guide*, SH26-4095
  - *IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server Installation and User’s Guide Version 5.3.2*, SC32-9059
  - *IBM Tivoli Storage Manager for Databases 5.5.1: Data Protection for Oracle for UNIX and Linux Installation and User’s Guide*, SC32-9064
  - *IBM Tivoli Storage Manager for Databases 5.5.1: Data Protection for Oracle for Windows Installation and User’s Guide*, SC32-9065
  - *IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Installation and User’s Guide for DB2 Version 5.5.0*, SC33-6341
  - *IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Installation and User’s Guide for Oracle V5.5*, SC33-6340
  - *IBM Tivoli Storage Manager for Hardware: Data Protection for Enterprise Storage Server for Databases (DB2 UDB) Installation and User’s Guide Version 5.2*, SC32-9060
  - *IBM Tivoli Storage Manager for Hardware: Data Protection for Enterprise Storage Server Databases (Oracle) Installation and User’s Guide Version 5.2*, GC32-1772
  - *IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino for UNIX, Linux, and OS/400 Installation and User’s Guide*, SC32-9056
  - *IBM Tivoli Storage Manager for Mail 5.5.1: Data Protection for Lotus Domino for Windows Installation and User’s Guide*, SC32-9057
  - *IBM Tivoli Storage Manager for Mail 5.5.1: Data Protection for Microsoft Exchange Server Installation and User's Guide Version 5.5.0*, SC32-9058

Online resources

These Web sites are also relevant as further information sources:

- IBM Tivoli Storage Manager documentation

  http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp

- IBM TotalStorage Productivity Center for Data

  http://www.ibm.com/servers/storage/software/center/data/index.html

- IBM Tivoli Storage Manager Express

  http://www.ibm.com/software/tivoli/products/storage-mgr-express/
- IBM Tivoli Storage Manager information

  http://www.ibm.com/software/tivoli/products/storage-mgr/
- IBM Tivoli Storage Manager Extended Edition information

  http://www.ibm.com/software/tivoli/products/storage-mgr-extended/
- IBM Training

  http://www.ibm.com/software/sw-training/
- IBM Virtual Innovation Center

  http://www.ibm.com/partnerworld/vic
- IBM Learning Services

  http://www.ibm.com/services/learning/index.html
- IBM Software Support Web page

  http://www.ibm.com/software/support/index_A_Z.html
- IBM FTP Software Web page:

  http://ftp.software.ibm.com/storage/tivoli-storage-management/

How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

T

Deployment Guide Series: IBM Tivoli Storage Manager V5.5

(0.5” spine)
0.475”<->0.875”
250 <-> 459 pages

Deployment Guide Series: IBM Tivoli Storage Manager V5.5

Plan for IBM Tivoli Storage Manager deployment

Perform installation and configuration

Learn from example scenarios

The objective of this IBM Redbooks publication is to provide comprehensive instructions for deploying the IBM Tivoli Storage Manager to various environments.

Readers should possess general knowledge of communication network architecture and design, basic size considerations of the Tivoli Storage Manager database, and basic pool management of Tivoli Storage Manager servers. This document is intended to be read and used by presales systems engineers and services personnel to build a customized deployment of the Tivoli Storage Manager. We expect such individuals have significant knowledge of Tivoli Storage Manager, and ideally the reader has attended Tivoli Storage Manager basic and advanced training classes.

The reader must be familiar with the following topics:

- Storage management concepts
- Network topologies
- Distributed systems architectures and configuration

This book is a valuable addition to, and can be read in conjunction with the existing product documentation.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-7379-01 ISBN 0738431672